University of Missouri

VIMAN Lab
VIRTUALIZATION, MULTIMEDIA AND NETWORKING LAB

AWS Concepts and Lab Intro

Saptarshi Debroy, Minh Nguyen

Contact: Prasad Calyam

Slides adapted from Hwang, Fox, Dongarra
&
Programming Amazon EC2, Vliet, Paganelli

Growth of AWS

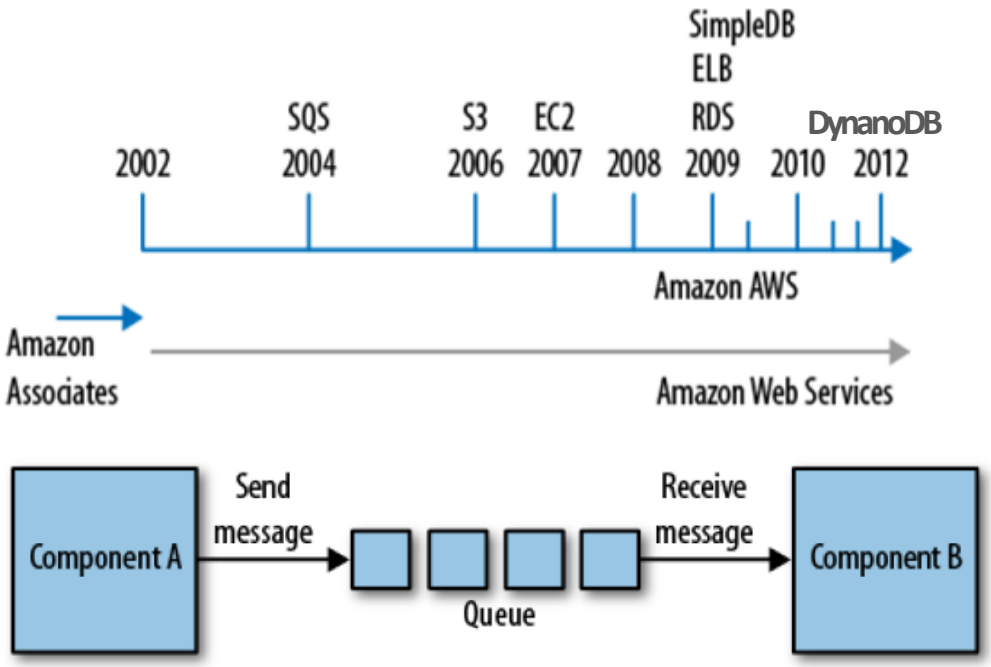

1.3 Trillion in 2013

Scale: Amazon S3 Growth

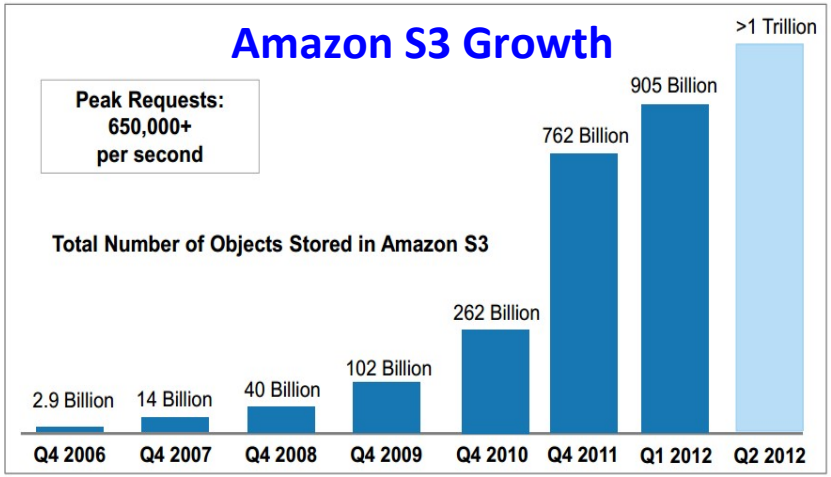

- AWS “Simple Queue Service” fosters ‘decoupled’ service oriented architecture message passing
  - “Developers can move data between distributed components of their applications that perform different tasks, without loosing messages or requiring each component to be always available”
- S3 Storage, EC2 Compute, ELB Load Balancing, RDS Relational Database, SimpleDB

AWS Free-Tier

Free Usage Restrictions

AWS Free Tier (Per Month):

Elastic Compute Cloud (EC2)

- 750 hours of Amazon EC2 Linux t.2 micro instance usage (1 GiB of memory and 32-bit and 64-bit platform support) – enough hours to run continuously each month*
- 750 hours of Amazon EC2 Microsoft Windows Server† t.2 micro instance usage (1 GiB of memory and 32-bit and 64-bit platform support) – enough hours to run continuously each month*
- 750 hours of an Elastic Load Balancer plus 15 GB data processing*
- 30 GB of Amazon Elastic Block Storage in any combination of General Purpose (SSD) or Magnetic, plus 2 million I/Os (with EBS Magnetic) and 1 GB of snapshot storage*

Simple Storage Service (S3)

- 5 GB of Amazon S3 standard storage, 20,000 Get Requests, and 2,000 Put Requests*

DynamoDB

25 GB of Storage, 25 Units of Read Capacity and 25 Units of Write Capacity – Enough to handle up to 200M requests per month with Amazon DynamoDB.**

Relational Database Service (RDS)

- 750 hours of Amazon RDS Single-AZ Micro DB Instances, for running MySQL, PostgreSQL, Oracle BYOL or SQL Server (running SQL Server Express Edition) – enough hours to run a DB Instance continuously each month*
- 20 GB of database storage, in any combination of RDS General Purpose (SSD) or Magnetic storage
- 10 million I/Os (for use with RDS Magnetic storage; I/Os on RDS General Purpose (SSD) storage are not separately billed)
- 20 GB of backup storage for your automated database backups and any user-initiated DB Snapshots

Simple Workflow (SWF)

- 1,000 Amazon SWF workflow executions and a total of 10,000 activity tasks, signals, timers and markers, and 30,000 workflow-days.**

Simple Queue Service (SQS) and Simple Notification Service (SNS)

- 1,000,000 Requests of Amazon Simple Queue Service**
- 1,000,000 Requests, 100,000 HTTP notifications and 1,000 email notifications for Amazon Simple Notification Service**

Amazon Elastic Transcoder

- 20 minutes of SD transcoding or 10 minutes of HD transcoding**

CloudWatch

- 10 Amazon Cloudwatch metrics, 10 alarms, and 1,000,000 API requests**

Data Transfer

- 15 GB of bandwidth out aggregated across all AWS services*

Data Pipeline

- 3 low frequency preconditions running on AWS per month*
- 5 low frequency activities running on AWS per month*

ElastiCache

- 750 hours of Amazon ElastiCache - enough hours to run a Cache Node continuously each month.*

Amazon Mobile Analytics

- 100 million free events per month**

$100 Promotional Code/Student Tier

You can access -

- AmazonRedshift
- **AWSDirectConnect**
- AmazonCloudcast
- AWSQueueService
- **AmazonVPC**
- AmazonElastiCache
- AmazonSES
- AmazonSIS
- AmazonCloudSearch
- AmazonSNS
- AmazonRoute53
- AWSStorageGateway
- **AmazonEC2**
- AmazonDynamoDB
- **ElasticMapReduce**
- Amazon ETS
- AmazonSimpleDB
- **AmazonRDS**
- AWSDataTransfer
- AWSSupportBasic
- **AmazonS3**
- **AmazonCloudFront**
- AWSElasticBeanstalk
- AmazonGlacier
- AWS Lambda
- AWS Key Management Service
- **CloudWatch**
- AWS cognito
- **Amazon ELB**

- **Students will receive a $100 AWS usage credit code from me/TA via e-mail**
- **Although each assigned lab session will only use free-tier resources**, the credit is helpful if there are accidental charges or if a student would like to experiment with any advanced AWS capabilities
- **If a student exceeds the $100 usage credit, he/she will be responsible for payment of any overage charges**

Example Application Hosting in AWS

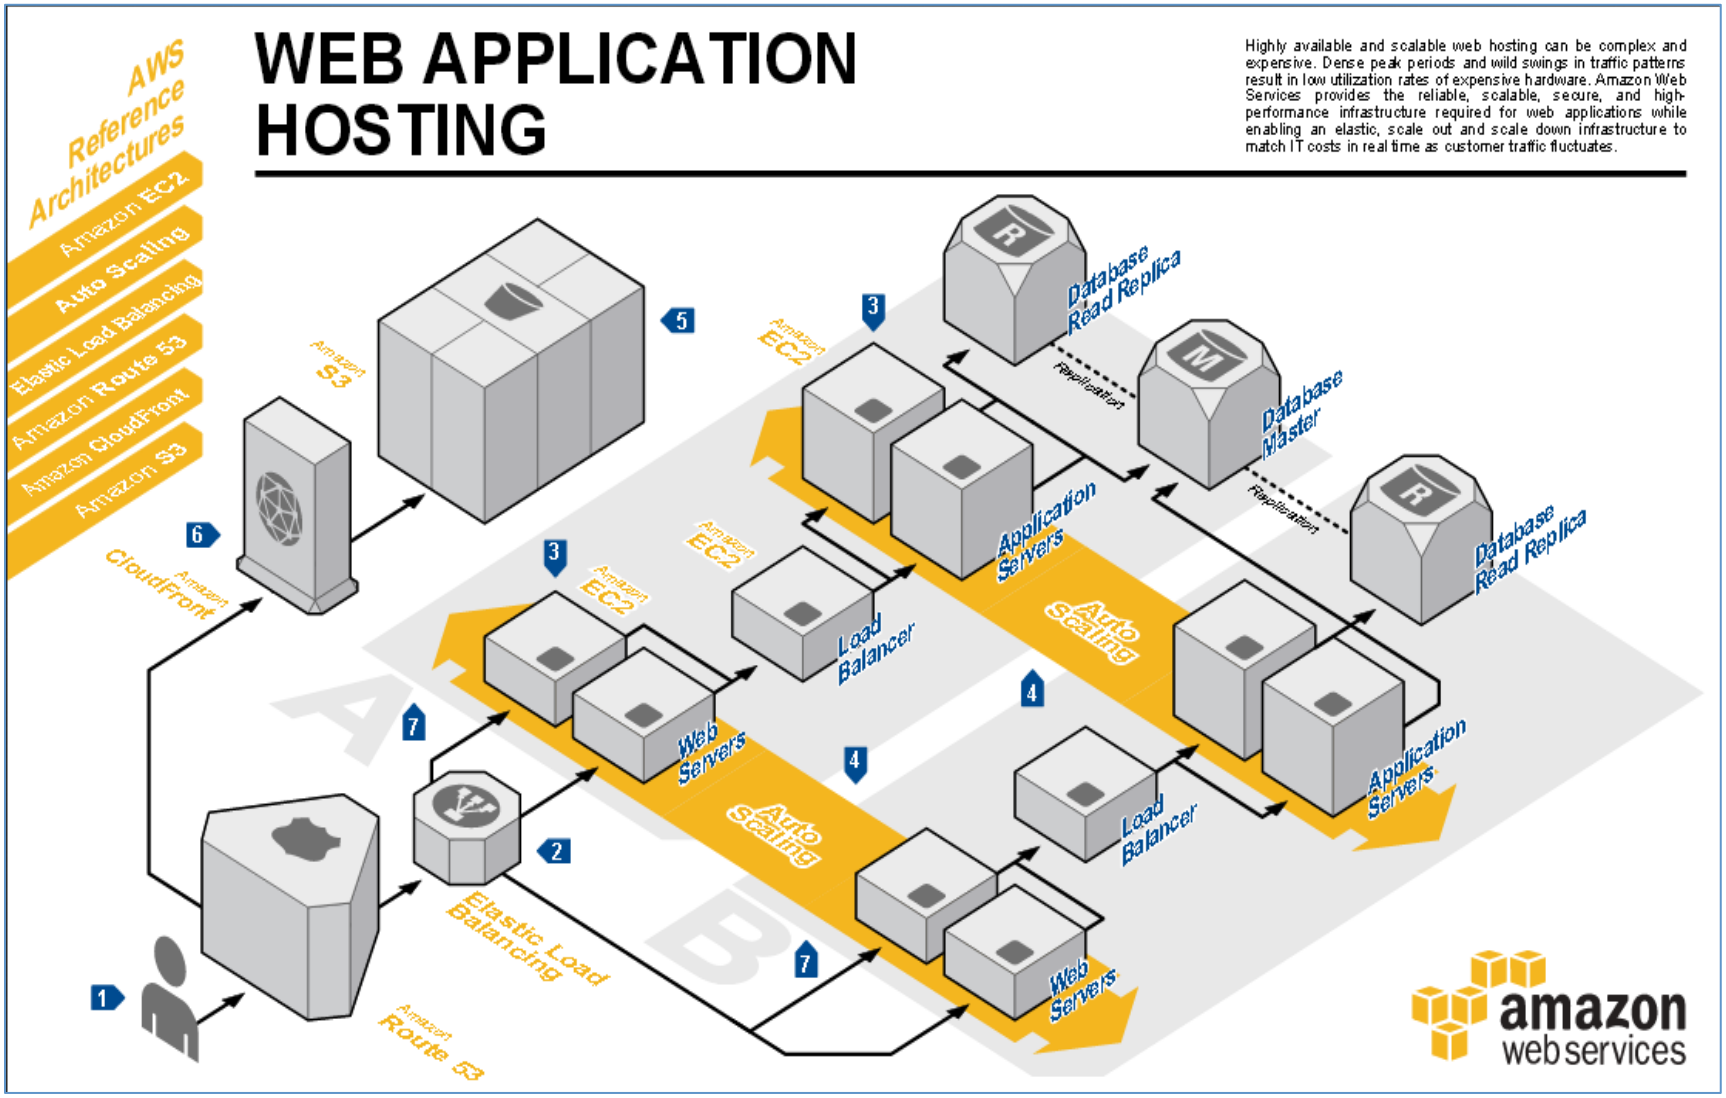

AWS Regions

11 regions, 28 availability zones (1 to 6 data centers)…
~1.4 million servers worldwide!!

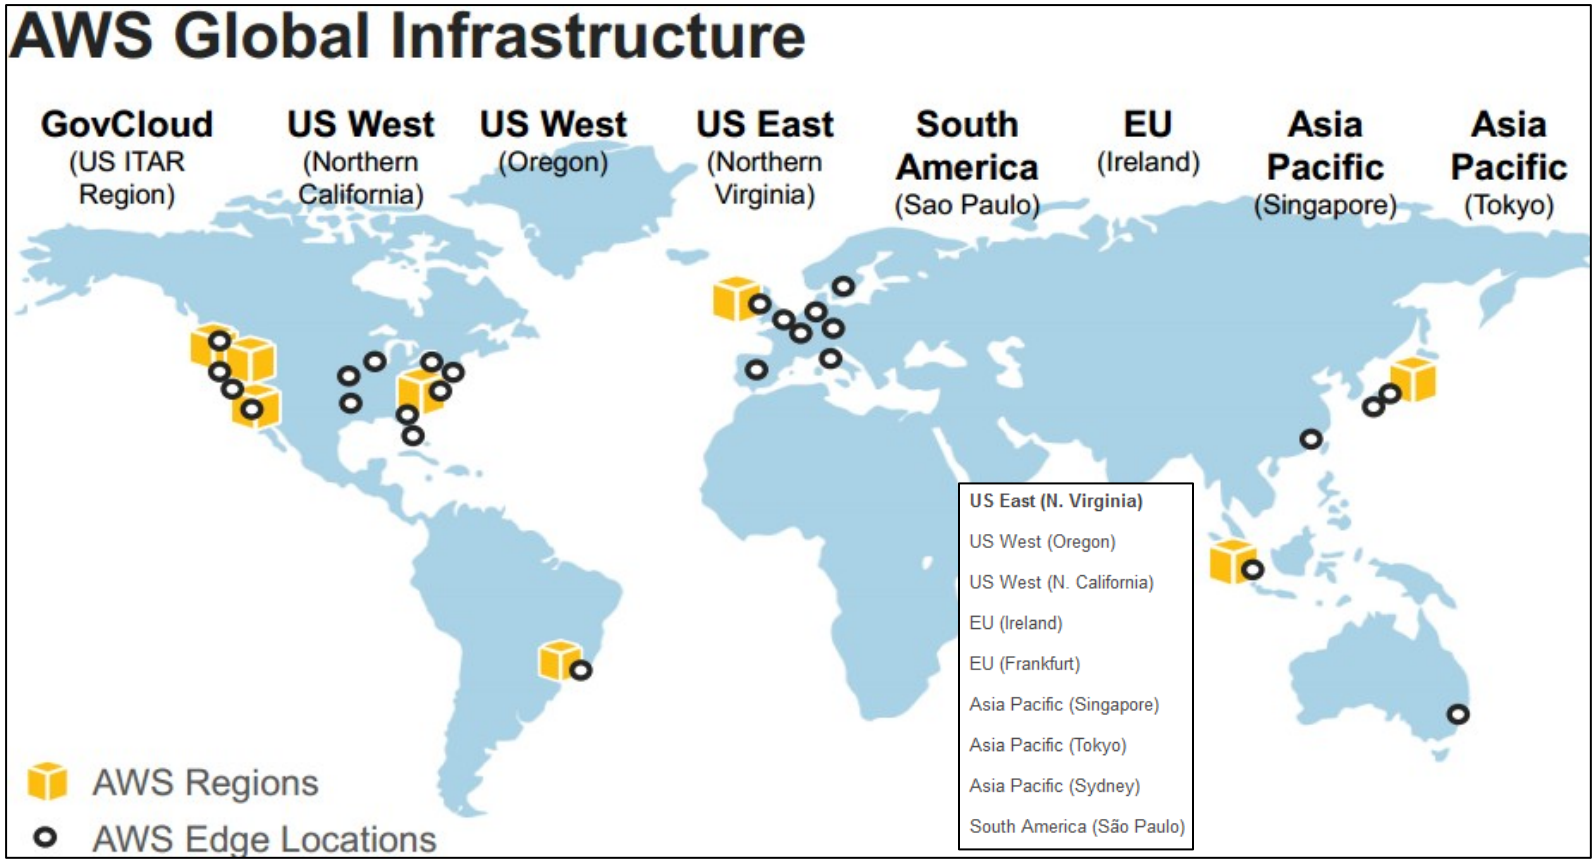

More Introduction Information about AWS at – (especially see different networking setups allowed by AWS including 'Direct Connect')
http://www.slideshare.net/AmazonWebServices/03-introduction-to-aws

Interesting must-read article…

- T. Morgan, “A Rare Peek Into The Massive Scale of AWS”, Nov. 2014
  - http://www.enterprisetech.com/2014/11/14/rare-peek-massive-scale-aws/

Lab Steps

- **Lab – AWS Account Setup, Services Overview, Resource Discovery, and Instance Setup**
- **Purpose of the Lab**
  - Understand definitions of various Amazon Web Services (AWS) and their use in cloud computing based web applications that are accessible over the Internet through an AWS account
  - Use the AWS account for the discovery, reservation and access of virtual compute/storage infrastructure instances

AWS Platform Example Deployment

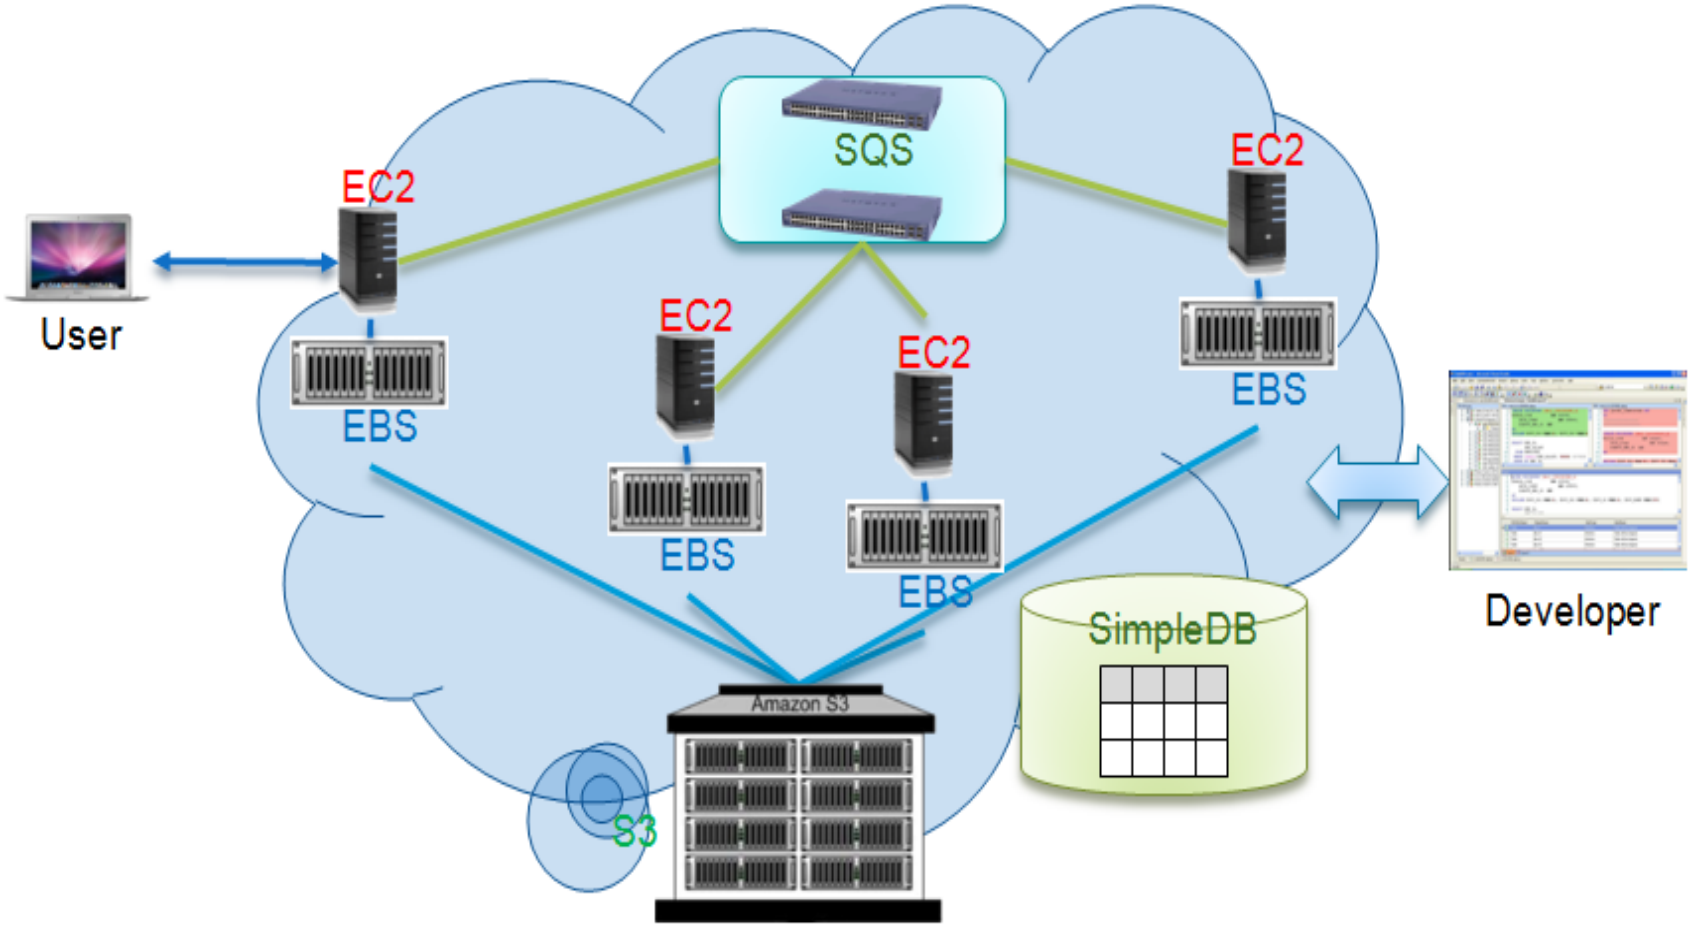

AWS Execution Environment

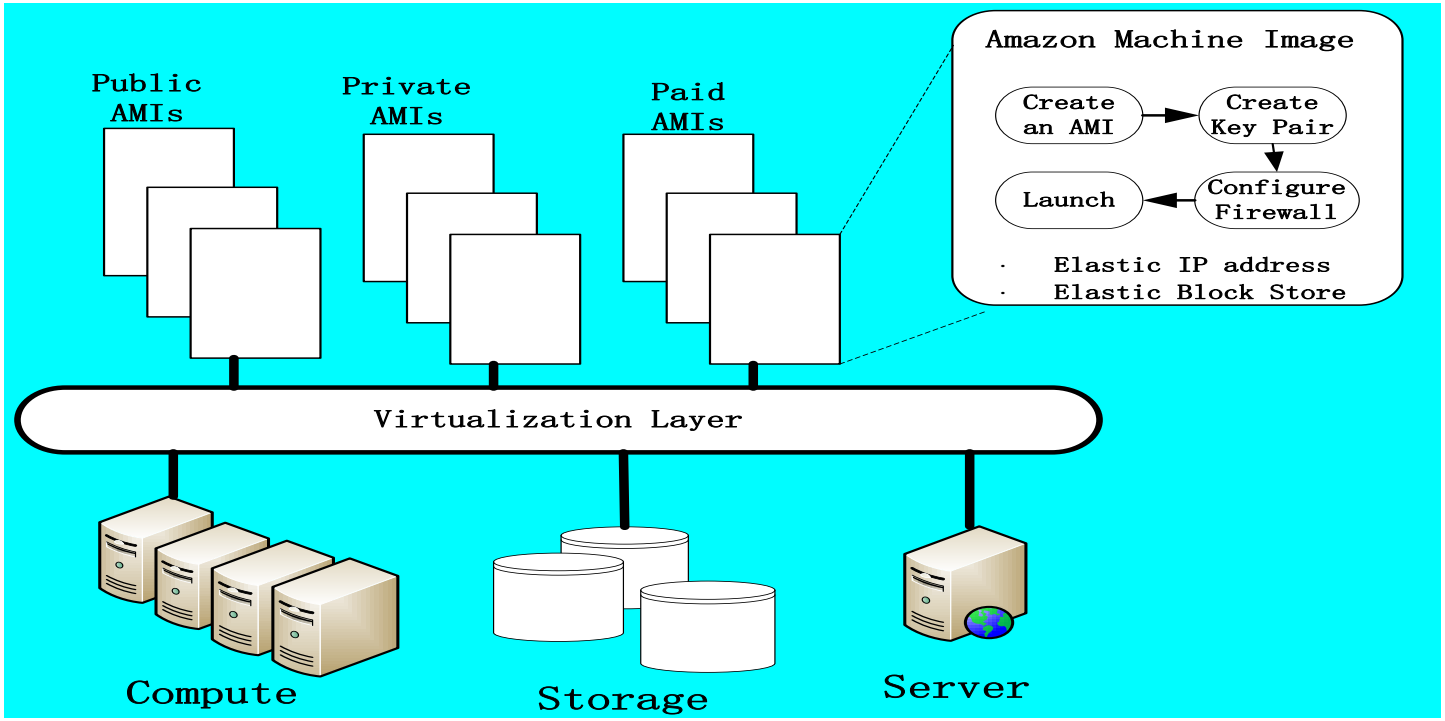

- **Private AMI:** Images created by you, which are private by default; you can grant access to other users to launch your private images
- **Private AMI:** Images created by users and released to the community, so anyone can launch instances based on them and use them any way they like
- **Paid AMI:** You can create images providing specific functions that can be launched by anyone willing to pay you per each hour of usage on top of AWS charges

AWS Access Credentials

- Credential type you use depends on the type of AWS API
  - **Access Keys**
    - To make secure REST or Query protocol requests to any AWS service API
    - Parts and Usage
      - *Access Key ID*—Your Access Key ID identifies you as the party responsible for service requests; you include it in each request, so it's not a secret
      - *Secret Access Key*—Each Access Key ID has a Secret Access Key associated with it; This key is used to calculate the digital signature that you include in the request; Your Secret Access Key is a secret, and only you and AWS should have it
  - **X.509 Certificates**
    - To make secure SOAP protocol requests to AWS service APIs
    - Parts and Usage
      - *X.509 Certificate* – holds the public key and related metadata; You include it in each service request, so it's not a secret
      - *Private Key*—Each certificate has a private key associated with it; Use the private key to calculate the digital signature to include in the request; Your private key is a secret, and only you should have it and AWS doesn't keep a copy
  - **Key Pairs**
    - To launch and then securely access your Amazon EC2 instances
    - You can make as many as you like by giving friendly names (can't replace any particular key pair)
    - Private key that you keep with you; Public key that AWS keeps to allow access

Install your first AWS-hosted Web App!

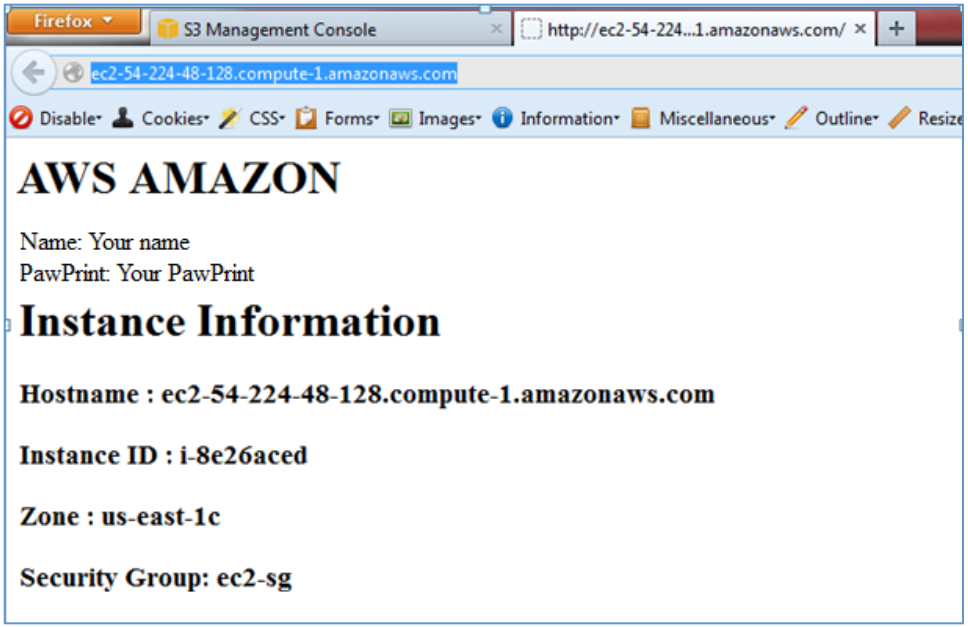

- Install HTTP, PHP and MySQL – LAMP package in your instance
- Read Hostname, Instance ID, Zone and Security Group from Instances Data from metadata set and show it on the web page
- Clean-up resources – remove snapshot, detach/remove volume

Recap: Example Application Hosting in AWS

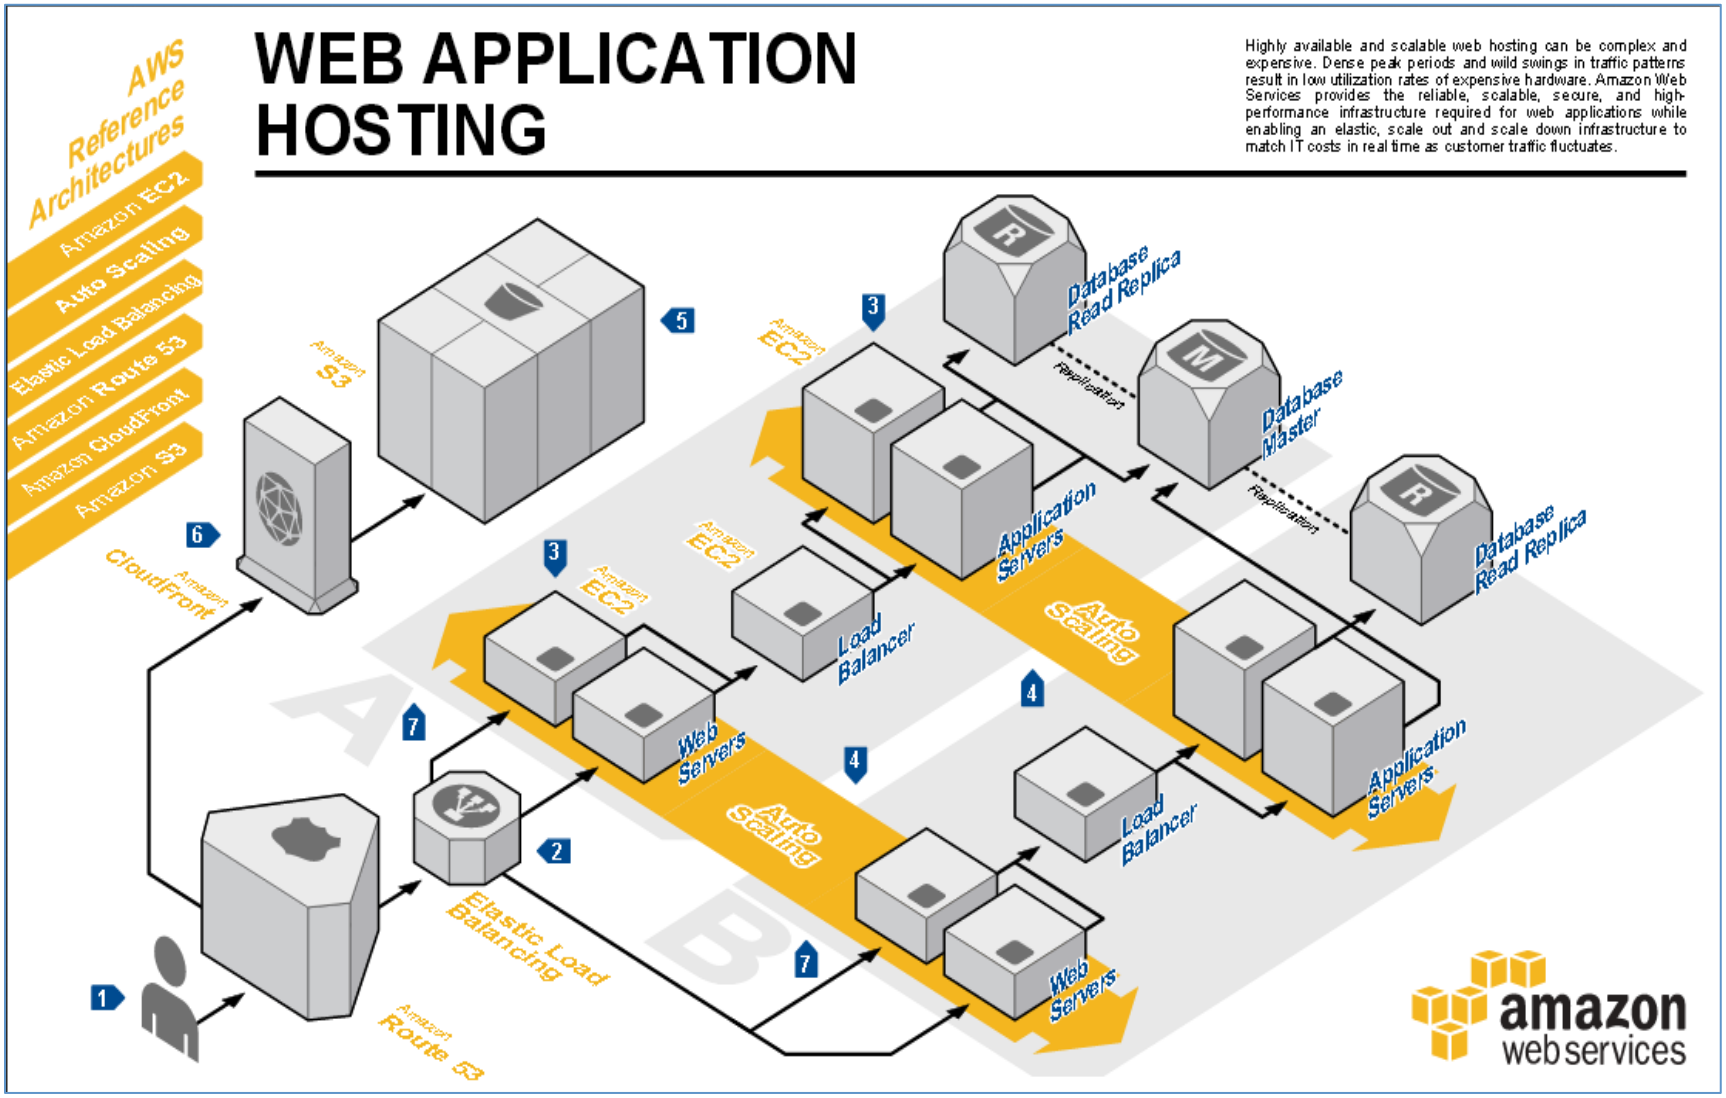

Cost Saving Considerations in AWS

- **On-Demand Instances**
  - Pay for compute capacity by the hour with no long-term commitments
- **Reserved Instances**
  - Make a low, one-time payment for each instance you want to reserve and in turn receive a significant discount on the hourly charge for that instance
- **Spot Instances**
  - Bid on unused EC2 capacity and run those instances for as long as their bid exceeds the current Spot Price

Other Best Practices…

- AWS Lab Reading
  - http://media.amazonwebservices.com/AWS_Cloud_Best_Practices.pdf
    - Design for failure and nothing will fail
    - Decouple your components
    - Implement elasticity
    - Think parallel
    - Keep dynamic data closer to compute and static data closer to user
    - Know security and performance tradeoffs
- Another great link for high scalability, architecture case studies
  - http://highscalability.com - 'Building bigger, faster, more reliable websites'
  - YouTube Architecture
    - http://highscalability.com/youtube-architecture
  - Good Dashboard Example:
    - http://stackexchange.com/performance