

# Principles of Computer Science II

## Cloud Computing

Ioannis Chatzigiannakis

Sapienza University of Rome

### Lecture 3




## Virtualization

- ▶ Virtualization deals with “extending or replacing an existing interface so as to mimic the behavior of another system”
- ▶ Virtual system examples:
  - ▶ virtual private network,
  - ▶ virtual memory,
  - ▶ virtual machine,
  - ▶ ...




## Starting Point: Physical System

- ▶ Physical Hardware
  - ▶ Processors, Memory, I/O devices, ...
  - ▶ Physical resources often underutilized
  - ▶ Periods that are over-utilized
- ▶ Software:
  - ▶ Tightly coupled to Hardware,
  - ▶ Single active OS,
  - ▶ OS controls Hardware



## What is a Virtual Machine?

- ▶ Hardware-level Abstraction
  - ▶ Virtual Hardware: Processors, Memory, I/O devices, ...
  - ▶ Encapsulates all OS and application state.
- ▶ Virtualization Software:
  - ▶ Extra level of indirection decouples hardware and OS,
  - ▶ Multiplexes physical hardware across multiple “guest” VMs,
  - ▶ Strong isolation between VMs,
  - ▶ Manages physical resources, improves utilization.




## Virtual Machine Isolation

- ▶ Secure Multiplexing:
  - ▶ Run multiple VMs on single physical host,
  - ▶ Processor hardware isolates VMs.
- ▶ Strong Guarantees:
  - ▶ Software bugs, crashes, viruses within one VM cannot affect other VMs
- ▶ Performance Isolation:
  - ▶ Partition system resources,
  - ▶ Example: VirtualBox controls for reservation, limit, shares.




## Virtual Machine Encapsulation

- ▶ Entire VM in a file:
  - ▶ OS, applications, data;
  - ▶ Memory and device state.
- ▶ Snapshots and Clones:
  - ▶ Capture VM state on the fly and restore to point-in-time,
  - ▶ Rapid system provisioning, backup, remote mirroring.
- ▶ Easy Content Distribution:
  - ▶ Pre-configured apps, demos.
  - ▶ Virtual Appliances.



## Virtual Machine Compatibility

- ▶ Hardware Independent:
  - ▶ Physical hardware hidden by virtualization layer,
  - ▶ Standard virtual hardware exposed to VM.
- ▶ Create Once, Run Anywhere:
  - ▶ No configuration issues,
  - ▶ Migrate VMs between hosts.
- ▶ Legacy Virtual Machines:
  - ▶ Run legacy OS on new platform.



## Common Uses

- ▶ Test and Development
  - ▶ Rapidly provision test and development servers.
  - ▶ Store libraries of pre-configured test machines.
- ▶ Business Continuity
  - ▶ Reduce cost and complexity by encapsulating entire systems into single files
  - ▶ Replicated and restored on demand into any target system.
- ▶ Enterprise Desktop
  - ▶ Secure unmanaged PCs without compromising end-user autonomy by layering a security policy in software around desktop virtual machines.




## Common Uses

- ▶ Run legacy software on non-legacy hardware
- ▶ Run multiple operating systems on the same hardware
- ▶ Create a manageable upgrade path
- ▶ Manage outages (expected and unexpected) dynamically



## Virtualized Data Centers



Reduce costs by consolidating services onto the fewest number of physical machines



## Non-virtualized Data Centers

- ▶ Too many servers for too little work
- ▶ High costs and infrastructure needs
  - ▶ Maintenance
  - ▶ Networking
  - ▶ Floor space
  - ▶ Cooling
  - ▶ Power
  - ▶ Disaster Recovery




## Dynamic Data Centers

- ▶ Virtualization helps us break the “one service per server” model
- ▶ Consolidate many services into a fewer number of machines when workload is low, reducing costs
- ▶ Conversely, as demand for a particular service increases, we can shift more virtual machines to run that service
- ▶ We can build a data center with fewer total resources, since resources are used as needed instead of being dedicated to single services



# Towards Serverless Computing



# Function as a Service


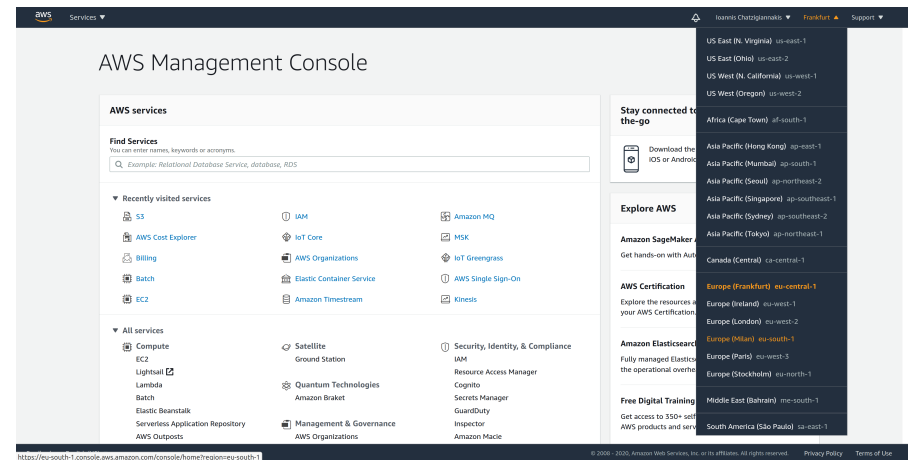
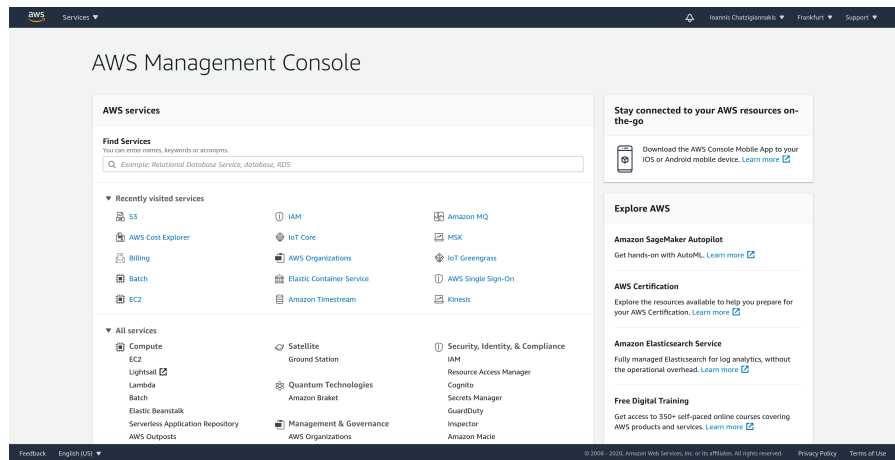


Figure 1. Magic Quadrant for Cloud Infrastructure and Platform Services

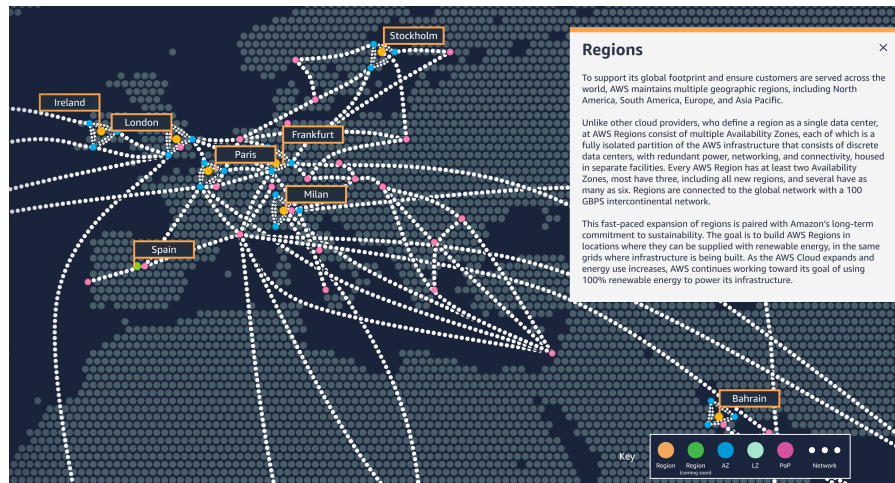


<ul style="list-style-type: none"> <li>▼ All services</li> <li>Computer <ul style="list-style-type: none"> <li>EC2</li> <li>LightSail</li> <li>Lambda</li> <li>Batch</li> <li>Elastic Beanstalk</li> <li>Serverless Application Repository</li> <li>AWS Outposts</li> <li>EC2 Image Builder</li> </ul> </li> <li>Containers <ul style="list-style-type: none"> <li>ECR</li> <li>Elastic Container Service</li> <li>Elastic Kubernetes Service</li> </ul> </li> <li>Storage <ul style="list-style-type: none"> <li>S3</li> <li>EFS</li> <li>Fsx</li> <li>S3 Glacier</li> <li>Storage Gateway</li> <li>AWS Backup</li> </ul> </li> <li>Database <ul style="list-style-type: none"> <li>RDS</li> <li>DynamoDB</li> <li>ElastiCache</li> <li>Neptune</li> <li>Amazon QLDB</li> <li>Amazon DocumentDB</li> <li>Amazon Keyspaces</li> <li>Amazon Timestream</li> </ul> </li> <li>Migration &amp; Transfer <ul style="list-style-type: none"> <li>AWS Migration Hub</li> <li>Application Discovery Service</li> <li>Database Migration Service</li> <li>Server Migration Service</li> <li>AWS AppConf</li> <li>AWS Transfer Family</li> <li>AWS Snow Family</li> <li>DataSync</li> </ul> </li> <li>Networking &amp; Content Delivery <ul style="list-style-type: none"> <li>VPC</li> <li>CloudFront</li> <li>Route 53</li> <li>API Gateway</li> <li>Direct Connect</li> <li>AWS App Mesh</li> <li>AWS Cloud Map</li> <li>Global Accelerator</li> </ul> </li> </ul>	<ul style="list-style-type: none"> <li>Developer Tools <ul style="list-style-type: none"> <li>CodeStar</li> <li>CodeCommit</li> <li>CodeArtifact</li> <li>CodeBuild</li> <li>CodeDeploy</li> <li>CodePipeline</li> <li>CodePipelines</li> <li>Cloud9</li> <li>X-Ray</li> </ul> </li> <li>Customer Enablement <ul style="list-style-type: none"> <li>AWS IQ</li> <li>SageMaker</li> <li>Managed Services</li> <li>Activate for Startups</li> </ul> </li> <li>Robotics <ul style="list-style-type: none"> <li>AWS RoboMaker</li> </ul> </li> <li>Blockchain <ul style="list-style-type: none"> <li>Amazon Managed Blockchain</li> </ul> </li> <li>Satellite <ul style="list-style-type: none"> <li>Ground Station</li> </ul> </li> <li>Quantum Technologies <ul style="list-style-type: none"> <li>Amazon Braket</li> </ul> </li> <li>Management &amp; Governance <ul style="list-style-type: none"> <li>AWS Organizations</li> <li>CloudWatch</li> <li>AWS X-Ray</li> <li>CloudFormation</li> <li>CloudTrail</li> <li>Config</li> <li>OpsWorks</li> <li>Service Catalog</li> <li>Systems Manager</li> <li>AWS AppConfig</li> <li>Trusted Advisor</li> <li>Control Tower</li> <li>AWS License Manager</li> <li>AWS Well-Architected Tool</li> <li>Personal Health Dashboard</li> <li>AWS Chatbot</li> <li>Lambda Wizard</li> <li>AWS Compute Optimizer</li> <li>Resource Groups &amp; Tag Editor</li> </ul> </li> <li>Media Services <ul style="list-style-type: none"> <li>Kinesis Video Streams</li> <li>MediaConnect</li> <li>MediaConvert</li> <li>MediaLive</li> <li>MediaPackage</li> <li>MediaStore</li> <li>MediaTailor</li> <li>Elemental Appliances &amp; Software</li> <li>Amazon Interactive Video Service</li> <li>Elastic Transcoder</li> </ul> </li> </ul>	<ul style="list-style-type: none"> <li>Machine Learning <ul style="list-style-type: none"> <li>Amazon SageMaker</li> <li>Amazon Augmented AI</li> <li>Amazon CodeGuru</li> <li>Amazon Comprehend</li> <li>Amazon Forecast</li> <li>Amazon Fraud Detector</li> <li>Amazon Kendra</li> <li>Amazon Lex</li> <li>Amazon Personalize</li> <li>Amazon Polly</li> <li>AWS IQ</li> <li>Amazon Rekognition</li> <li>Amazon TestOps</li> <li>Amazon Transcribe</li> <li>Amazon Translate</li> <li>AWS DeepComposer</li> <li>AWS DeepLens</li> <li>AWS DeepRacer</li> </ul> </li> <li>Analytics <ul style="list-style-type: none"> <li>Athena</li> <li>Amazon Redshift</li> <li>EMR</li> <li>CloudSearch</li> <li>Elasticsearch Service</li> <li>Kinesis</li> <li>QuickSight</li> <li>Data Pipeline</li> <li>AWS Data Exchange</li> <li>AWS Glue</li> <li>AWS Lake Formation</li> <li>MSK</li> </ul> </li> <li>Security, Identity, &amp; Compliance <ul style="list-style-type: none"> <li>IAM</li> <li>Resource Access Manager</li> <li>IOT Core</li> <li>Finegrained</li> <li>GuardDuty</li> <li>Inspector</li> <li>Amazon Macie</li> <li>AWS Single Sign-On</li> <li>Certificate Manager</li> <li>Key Management Service</li> <li>CloudHSM</li> <li>Directory Service</li> <li>USF &amp; Shield</li> <li>AWS Firewall Manager</li> <li>Artifact</li> <li>Security Hub</li> <li>Detective</li> </ul> </li> <li>AWS Cost Management <ul style="list-style-type: none"> <li>AWS Cost Explorer</li> <li>AWS Budgets</li> <li>AWS Marketplace Subscriptions</li> </ul> </li> </ul>	<ul style="list-style-type: none"> <li>Front-end Web &amp; Mobile <ul style="list-style-type: none"> <li>AWS Amplify</li> <li>Mobile Hub</li> <li>AWS AppSync</li> <li>Device Farm</li> </ul> </li> <li>AR &amp; VR <ul style="list-style-type: none"> <li>Amazon Sumerian</li> </ul> </li> <li>Application Integration <ul style="list-style-type: none"> <li>Step Functions</li> <li>Amazon AppFlow</li> <li>Amazon EventBridge</li> <li>Amazon MQ</li> <li>Simple Notification Service</li> <li>Simple Queue Service</li> <li>SWF</li> </ul> </li> <li>Customer Engagement <ul style="list-style-type: none"> <li>Amazon Connect</li> <li>Pinpoint</li> <li>Simple Email Service</li> </ul> </li> <li>Business Applications <ul style="list-style-type: none"> <li>Alexa for Business</li> <li>Amazon Chime</li> <li>WorkMail</li> <li>Amazon WorkDocs</li> </ul> </li> <li>End User Computing <ul style="list-style-type: none"> <li>WorkSpaces</li> <li>AppStream 2.0</li> <li>WorkSpaces</li> <li>WorkLink</li> </ul> </li> <li>Internet of Things <ul style="list-style-type: none"> <li>IOT Core</li> <li>FreeRTOS</li> <li>IOT 1-Click</li> <li>IOT Analytics</li> <li>IOT Device Defender</li> <li>IOT Device Management</li> <li>IOT Events</li> <li>IOT Greengrass</li> <li>IOT SiteWise</li> <li>IOT Things Graph</li> </ul> </li> <li>Game Development <ul style="list-style-type: none"> <li>Amazon GameLift</li> </ul> </li> </ul>
---	---	---	--

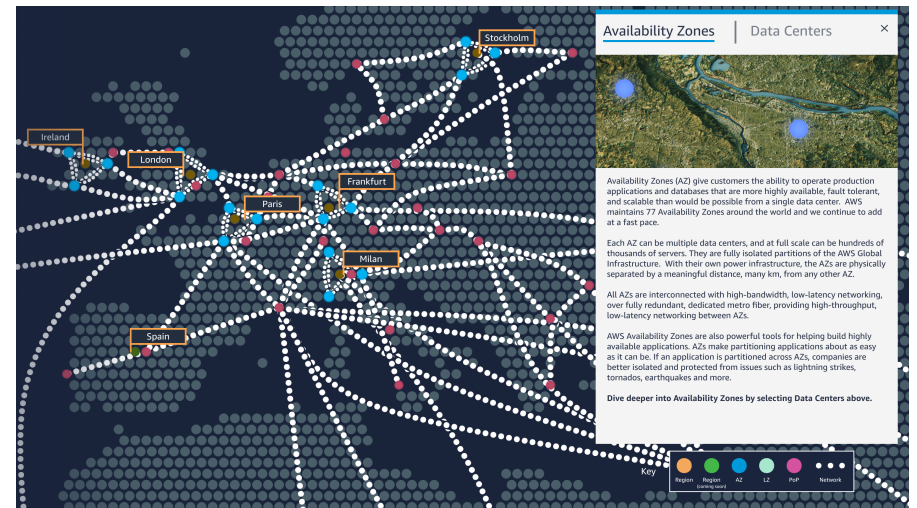




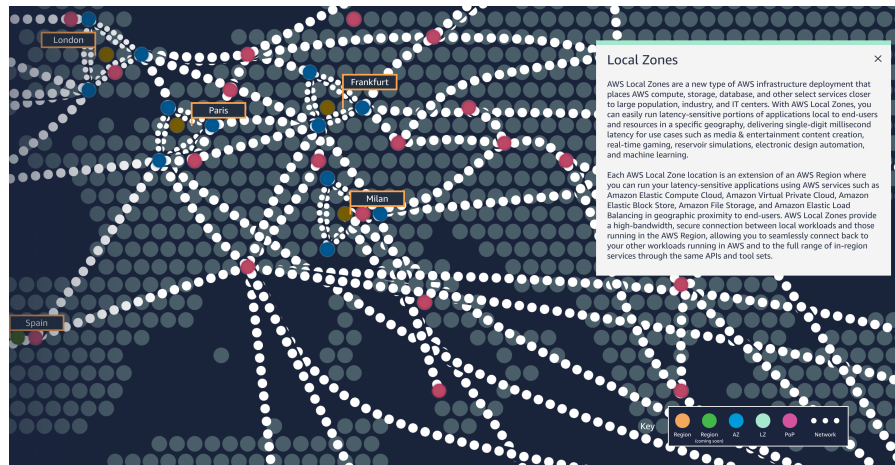
## AWS Infrastructure



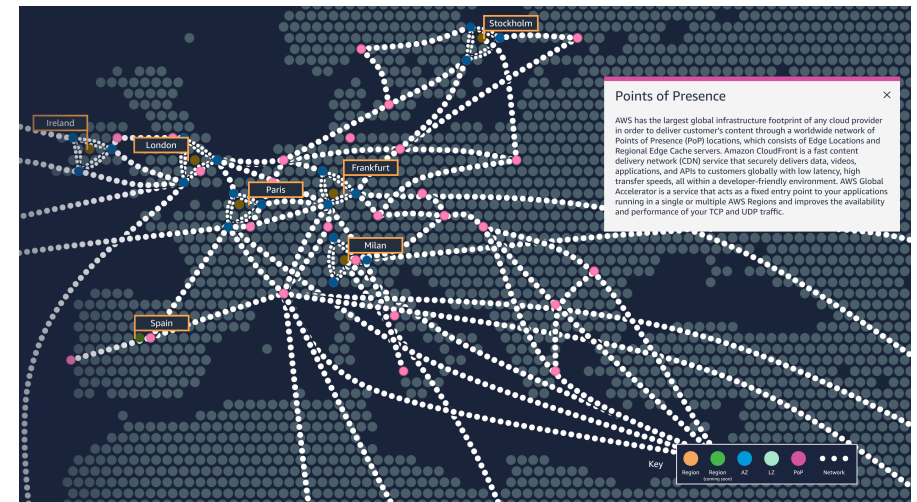
## AWS Infrastructure



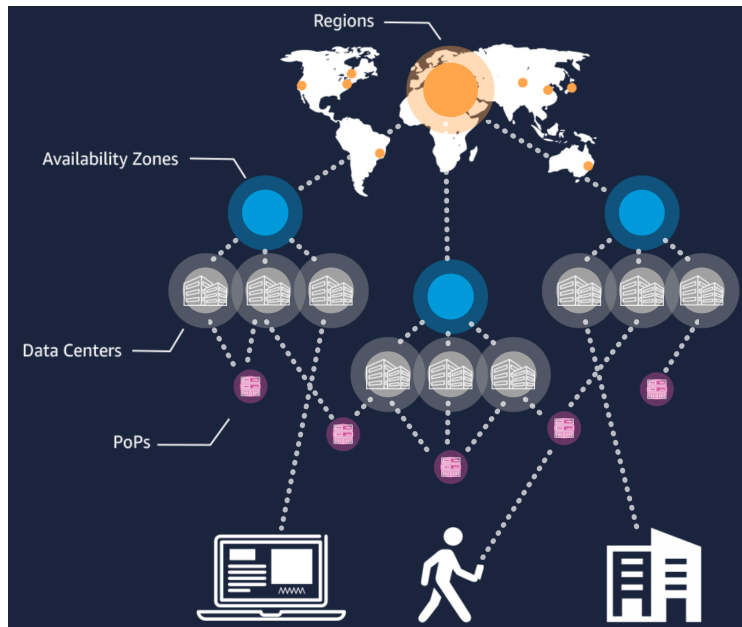
## AWS Infrastructure



## AWS Infrastructure



## AWS Infrastructure



## Introduction to AWS S3

- ▶ S3 = Simple Storage Service
  - ▶ From 0 bytes to 5 Tbytes.
- ▶ Provides a secure, durable, highly-scalable storage space.
  - ▶ AWS secures content with encryption, ACL and bucket policies.
  - ▶ AWS guarantees 99.99999999% durability (11 × 9s).
  - ▶ AWS guarantees 99.99% availability.
- ▶ We can access items stored:
  - ▶ Using the web.
  - ▶ Using the Web Console.
  - ▶ Using the Smartphone App.
  - ▶ From the Command line AWS tool.
  - ▶ Programmatically through the AWS S3 API.



## S3 Basics

- ▶ Object-based storage.
  - ▶ Files = Objects.
  - ▶ Not suitable to install an operating system or host a database.
- ▶ Files/Objects are organized in Buckets.
- ▶ Bucket names must be unique – **S3 is a universal namespace**.
  - ▶ <http://sapienza2020adm.s3.amazonaws.com/>
  - ▶ When you create a new S3 bucket, AWS creates a new web address.
- ▶ Objects (Files) have the following properties:
  - ▶ Key: the name of the object.
  - ▶ Value: the actual contents.
  - ▶ Version ID: used by the versioning system.
  - ▶ Metadata: tags that we can attach to objects.
  - ▶ ACL: who can access the object.



## S3 Storage Classes

- ▶ Free Tier – new AWS accounts
  - ▶ 5GB of S3 storage.
  - ▶ 20,000 GET – 2,000 PUT/COPY/POST/LIST
  - ▶ 15GB of Data Transfer Out each month for one year
- ▶ S3 Standard
  - ▶ \$0.0245 per GB
  - ▶ \$0.0054 per 1000 PUT/COPY/POST/LIST
  - ▶ \$0.00043 per 1000 GET/SELECT/all other requests.
- ▶ S3-IA Infrequent Access
  - ▶ \$0.0135 per GB – a minimum storage duration of 30 days.
  - ▶ \$0.01 per 1000 PUT/COPY/POST/LIST
  - ▶ \$0.001 per 1000 GET/SELECT/all other requests.
- ▶ S3 Glacier
  - ▶ \$0.0045 per GB – a minimum storage duration of 90 days.
  - ▶ \$0.06 per 1000 PUT/COPY/POST/LIST
  - ▶ \$0.00043 per 1000 GET/SELECT/all other requests.



## 2<sup>nd</sup> Assignment

- ▶ <https://www.rosalind.info/>
  - ▶ Complete the following **challenges**:  
dna, rna, revc, iprb, gc, subs, lia, iev, cons, prob
  - ▶ <http://rosalind.info/problems/{challenge}>
- ▶ Create a GitHub repository and upload the code for each exercise.
- ▶ Email [ichatz@diag.uniroma1.it](mailto:ichatz@diag.uniroma1.it)  
Subject: [PCS2] Homework 2  
Your GitHub repository with your solutions, for all challenges.  
Also send your account user account link:  
<http://rosalind.info/users/{username}>

