CLOUD COMPUTING

Introduction to Amazon Web Services Amazon Simple Storage Service - S3

Leonidas Akritidis – lakritidis@ihu.gr

Introduction to AWS

- Amazon Web Services is a complete set of cloud-based products and services offered by Amazon.
- AWS includes tools for distributed and fault tolerant:
 - Storage & databases,
 - Computational resources,
 - Networking & content delivery services,
 - (Big) Data analytics,
 - Machine Learning algorithms,
 - Web hosting platforms
 - and numerous (more than 175) others

AWS Infrastructure – Data Centers

- AWS is the largest cloud infrastructure in the world.
- It includes 22 Data Centers (DC) in 22 regions scattered across the planet.
- A data center is a gigantic installation of interconnected computers and networking equipment.

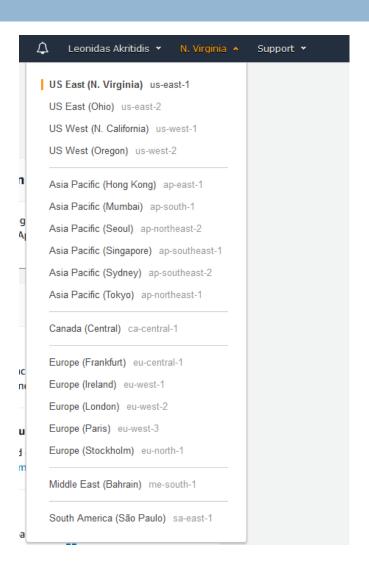


AWS Infrastructure – AZ

- Additional fault tolerance is achieved through the maintenance of Availability Zones (AZ).
- Each data center includes at least three such Availability Zones. These are isolated "areas" within the data center, where the data and critical operations are replicated.
- When an availability zone goes down (due to power loss, or hardware failure, etc.), another one replaces it immediately with its replica data (from the same DC).
- With this policy, AWS has managed to become the most stable cloud provider, with the smallest "down times".

AWS Infrastructure – AZ & Services

□ When the user selects a service from AWS, s/he is free to select the data center (region) where the service will be deployed.



History

aws

Console Home

Compute

EC2

Lightsail @

ECR

ECS

EKS

Lambda

Batch

Elastic Beanstalk

Serverless Application Repository

AWS Outposts

EC2 Image Builder

Storage

S3

EFS FSx

S3 Glacier

Storage Gateway

AWS Backup



Database

RDS

DynamoDB

ElastiCache

Neptune

Amazon Redshift

Amazon QLDB

Amazon DocumentDB

Managed Cassandra Service



Migration & Transfer

AWS Migration Hub

Application Discovery Service

Database Migration Service

Server Migration Service

AWS Transfer for SFTP



Customer Enablement

AWS IQ Z

Support

Managed Services



Amazon Managed Blockchain



Satellite

Ground Station



Quantum Technologies

Amazon Braket Z



Management & Governance

AWS Organizations

CloudWatch

AWS Auto Scaling

CloudFormation

CloudTrail

Config

OpsWorks

Service Catalog

Systems Manager

AWS AppConfig

Trusted Advisor

Control Tower

AWS License Manager

AWS Well-Architected Tool

Personal Health Dashboard 2

AWS Chatbot

Launch Wizard

AWS Compute Optimizer



Media Services



Machine Learning

Amazon SageMaker

Amazon CodeGuru

Amazon Comprehend

Amazon Forecast

Amazon Fraud Detector

Amazon Kendra

Amazon Lex

Amazon Machine Learning

Amazon Personalize

Amazon Pollv

Amazon Rekognition

Amazon Textract

Amazon Transcribe

Amazon Translate

AWS DeepLens

AWS DeepRacer

Amazon Augmented Al



Athena

EMR

CloudSearch

Elasticsearch Service

Kinesis

QuickSight 2

Data Pipeline

AWS Data Exchange

AWS Glue

AWS Lake Formation

MSK

Security, Identity, & Compliance

IAM

Resource Access Manager

Cognito

GuardDutv

Secrets Manager

Application Integration

Step Functions

Amazon EventBridge

Amazon MQ

Simple Notification Service

Simple Queue Service

SWF



AWS Cost Management

AWS Cost Explorer

AWS Budgets

AWS Marketplace Subscriptions



Customer Engagement

Amazon Connect

Pinpoint

Simple Email Service



Business Applications

Alexa for Business

Amazon Chime Z

WorkMail

End User Computing

WorkSpaces

AppStream 2.0

WorkDocs

WorkLink



Internet Of Things

IoT Core

FreeRTOS

IoT 1-Click

IoT Analytics

IoT Device Defender

I-T Device Messesses

Amazon Account and Free-Tier

- To make use of AWS, an Amazon account is required.
- During the account creation, the user fills his/her personal data including credit card information.
- As of June 2020, Amazon offers a free-tier program which allows the usage of a limited amount of AWS resources for a limited period of time without charge.

Create an AWS account

Email address
ිස්ත්රිස ්ට@ihu.gr
Password
•••••
Confirm password
•••••
AWS account name 1
CloudInstructor
Continue

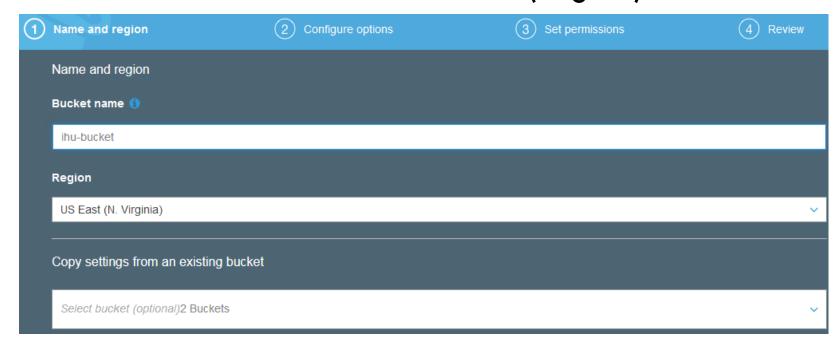
Sign in to an existing AWS account

Amazon S3 - Introduction

- S3 stands for Simple Storage Service.
- It is as simple as its name: a simple cloud-based storage framework designed for data availability (fault tolerance), scalability, security and performance.
- The user may upload files into S3 and access them from anywhere in the world.
- □ **Important:** The largest the files, the largest the amount AWS charges. These charges concern not only the file storage, but also, their transfer to S3 as well (for the allocation of AWS's network bandwidth).

Amazon S3 - Buckets

- □ The files in S3 are organized in buckets.
- A bucket is a fault-tolerant distributed storage space somewhere in the cloud.
 - □ The user determines the location (region) of the bucket.

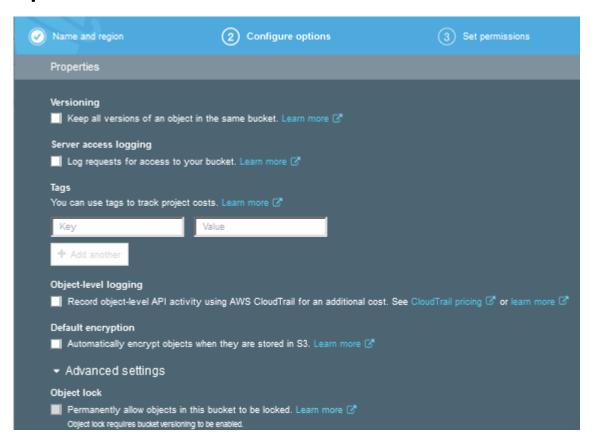


Amazon S3 - Bucket Features (1)

- Additional features for a S3 bucket include:
- Versioning.
 - Keep all versions of an object in the same bucket.
- Access Logging.
 - Logs all access requests to the bucket.
- Encryption.
 - Encrypts all the objects in the bucket for security.
- (Object-level) Locking.
 - Enables the locking of specific objects in the bucket.

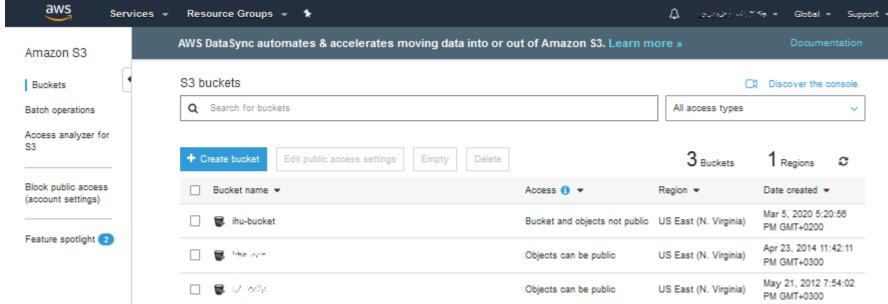
Amazon S3 – Bucket Features (2)

- Permissions
 - Allow public access to the contents of the bucket



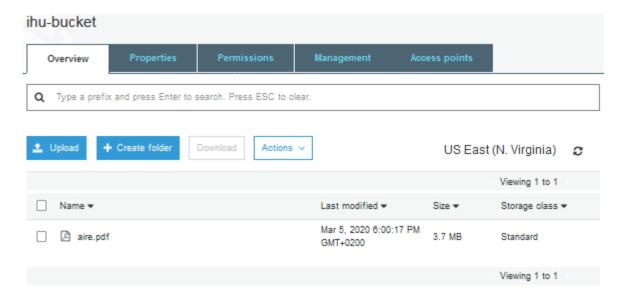
Amazon S3 – Bucket Management

- The recently created bucket appears in AWS S3 Management Console.
- From this point, new files can be appended into the bucket by clicking on its name.



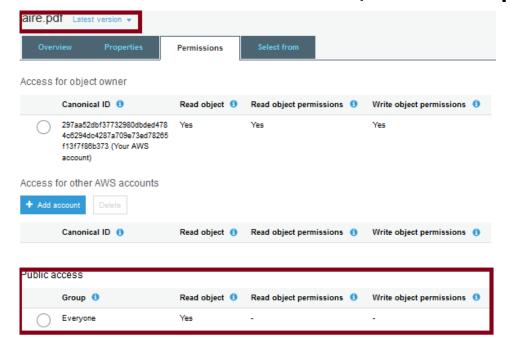
Amazon S3 — Bucket Management

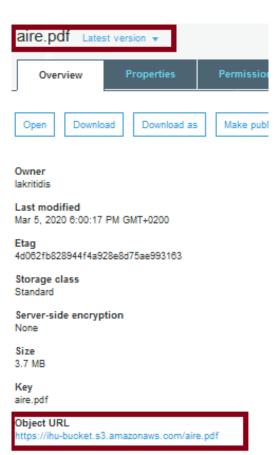
- From the 'Overview' Tab the bucket contents (files and folders) can be both managed and searched.
- □ In this example we uploaded 'aire.pdf' into 'ihu-bucket'.
- By clicking on this file AWS S3 presents the relevant information, details, permissions, etc.



Amazon S3 – Public Access

- To make a file publicly accessible, public permission must be granted. So click on the file, then on 'Permissions' tab.
- Grant public access to 'aire.pdf'.
- □ In the 'Overview' tab, see 'Object URL'.





Notes on Cloud Storage (1)

- AWS offers a convenient UI for managing the contents of a bucket in S3.
- It is similar to the conventional file managers of modern operating systems.
- In reality, the file organization and management in the cloud is very different than the 'local' operating systems.
- A bucket is just a resource in a data center of a region.
- No more information!

Notes on Cloud Storage (2)

- This could mean that the contents of a bucket may reside into:
 - the same hard drive, or
 - different hard drives on the same machine, or
 - different hard drives on different machines.
- Moreover, the same (large) file may span multiple hard drives (i.e. a part is in one drive and the other part/s are in different hard drive/s).
- Everything depends on the data distribution and partitioning
 framework/algorithm that runs on the cloud provider infrastructure.
- A considerable amount of modern research focuses on the design of effective/efficient data partitioning and distribution methods.
 - Requirements include load balancing, availability, fault tolerance, security, and performance.

More Resources

- □ Amazon S3 Getting Started
- Amazon S3 Developer Guide
- □ Working with Amazon S3 Buckets
- □ Working with Amazon S3 Objects