

# Introductory course to AWS Cloud and Core Services

## Elastic Cloud Compute (EC2)

Ioannis Chatzigiannakis

Sapienza University of Rome

### Lecture 2



## AWS: Elastic Compute Cloud (EC2)

- AWS EC2 = Elastic Compute Cloud
- Resizable compute resources in the cloud.
- Minimizes the time to provision a server.
  - Introduce a new server within minimum delay.
  - Scale capacity up very fast.
- Quickly modify the capabilities of the compute instance.
  - Introduce additional computational, memory and storage capabilities.
  - Reduce computational, memory and storage capabilities.
- Shutdown - or completely remove resources.
  - Scale down very fast.
- Pay only for the resources you need.



## Typical Use Cases

- Development and Testing Environments
- Hosting of Databases
- Hosting of web services
- Data analytics
- Code repository
- GPU-assisted machine learning
- High performance computing
- Video processing
- Backup and disaster recovery
- ...



## EC2 Provisioning Options

- **On Demand** – Pay for the compute capacity by the hour.
  - No up-front payment or long-term commitment.
  - Short-term, spiky, or unpredictable workloads.
  - Applications development or testing.
- **Spot Instances** – Acquire spare capacity up to 90% off the on-demand price.
  - When start/end times are flexible.
  - Applications that are only feasible at very low compute prices.
  - Urgent computing needs for large amounts of additional capacity.
- **Reserved Instances** – Significant discount (up to 75%) compared to On-Demand instance pricing.
  - For applications that have steady state or predictable usage.
  - Long term ( $\geq 1$  year) to reduce their total computing costs.
- **Dedicated Hosts** – Physical servers dedicated for use use.



## EC2 Instance Types

- **General Purpose** – balance of compute, memory and networking resources.
- **Compute Optimized** – ideal for compute bound applications that benefit from high performance processors.
- **Memory Optimized** – deliver fast performance for workloads that process large data sets in memory.
- **Accelerated computing** – use hardware accelerators, or co-processors, to perform functions, such as floating point number calculations, graphics processing, or data pattern matching, more efficiently than is possible in software running on generic CPUs.
- **Storage optimized** – for workloads that require high, sequential read and write access to very large data sets on local storage.



5 / 30

## EC2 Instance Types & Resources

- **CPU** – 64-bit Arm, AMD EPYC 7000, Intel Xeon Platinum 8175M, Intel Xeon E5-2676.
  - 1 ... 192 virtual CPUs – 1 thread = 1 vCPU.
- **Memory** – 1 ... 512 GB.
- **Network** – up to 100 Gbps.
- **Storage**
  - Amazon Elastic Block Store (EBS) – easy to use, high performance block storage service.
  - 0 ... 60 TB NVMe SSD – ensure best IOPS (Input/Output operations per second).
- **Hardware Accelerators**
  - NVIDIA Tesla V100 GPUs, NVIDIA K80 GPUs, NVIDIA T4 Tensor Core GPUs.
  - AWS Inferentia Chips.
  - Xilinx Virtex UltraScale+ VU9P FPGAs



6 / 30

## Available OS & Software

- **Operating Systems**
  - Linux/Unix – Amazon Linux, Debian, Ubuntu, Red Hat, CentOS, SUSE, FreeBSD, Gentoo, Mint, ...
  - Windows – Server 2019, Server 2016, Server 2012.
- **Databases** – PostgreSQL, MySQL, MongoDB, Neo4J, Oracle Enterprise, Microsoft SQL, ...
- **AWS Marketplace** – a wide selection of commercial and free software from well-known vendors.



7 / 30

## Pricing Examples

- **General Purpose**
  - **t2.micro Linux or Windows** – 2 vCPUs + 4 GB – 750 hours free per month, \$0.05/h
  - **a1.xlarge Linux** – 4 64-bit ARM vCPUs + 8 GB – \$0.1152/h
  - **a1.xlarge Linux** – 4 64-bit ARM vCPUs + 8 GB – \$0.1152/h
  - **m5.24xlarge Linux** – 96 Xeon vCPUs + 337 GB – \$5.136/h
  - **m5.24xlarge Windows** – 96 Xeon vCPUs + 337 GB – \$9.552/h
- **Compute Optimized**
  - **c5.xlarge Linux** – 4 Xeon vCPUs + 8 GB – \$0.192/Hour
  - **c5.24xlarge Linux** – 96 Xeon vCPUs + 192 GB – \$4.608/Hour
- **Hardware Accelerators**
  - **p3.2xlarge Linux** – 1 NVIDIA Tesla V100 GPUs + 8 Xeon vCPUs + 61 GB – \$3.305 per Hour
  - **p3dn.24xlarge Linux** – 8 NVIDIA Tesla V100 GPUs + 96 Xeon vCPUs + 768 GB – \$33.711 per Hour



8 / 30

## Amazon Elastic Block Store (EBS)

- Easy to use, high performance block storage service.
- Targeting both throughput and transaction intensive workloads.
  - Can be used for relational and non-relational databases.
  - Enterprise applications.
  - Big data analytics engines.
  - General purpose file systems.
  - Media workflows.
- Highly availability and durability – 99.999%
- Virtually unlimited scale – as little as a single GB of storage, or scale up to petabytes of data.
- Secure – encryption of data at-rest, data in-transit, and all volume backups.



## EBS Volume Types – HDD based

- **Throughput Optimized HDD (ST1)** – frequently accessed, throughput-intensive.
  - Large datasets and large I/O sizes, such as MapReduce, Kafka, log processing, data warehouse, and ETL workloads.
  - Low cost HDD volume.
  - Volume Size: 500 GB – 16 TB.
  - Max IOPS/Volume: 500
  - Max Throughput/Volume: 500 MB/s
  - Price: \$0.045/GB-month
- **Low-cost HDD (SC1)** – less frequently accessed large, cold datasets.
  - Colder data requiring fewer scans per day.
  - Volume Size: 500 GB – 16 TB.
  - Max IOPS/Volume: 250
  - Max Throughput/Volume: 250 MB/s
  - Price: \$0.025/GB-month

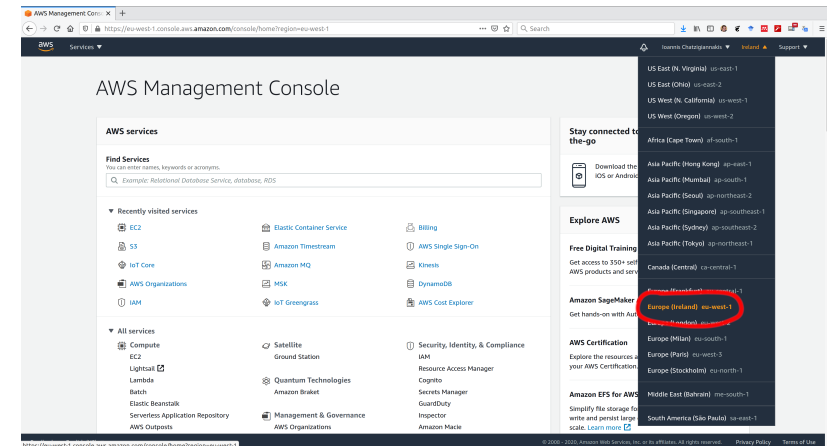


## EBS Volume Types – SSD based

- **Provisioned IOPS SSD (IO1)** – high performance SSD volume designed for latency-sensitive transactional workloads.
  - I/O-intensive NoSQL & relational databases.
  - Volume Size: 4 GB – 16 TB.
  - Max IOPS/Volume: 64,000
  - Max Throughput/Volume: 1,000 MB/s
  - Price: \$0.125/GB-month + \$0.065/provisioned IOPS
- **Default EBS volume type (GP2)** – ideal for suitable for a broad range of transactional workloads.
  - Boot volumes, low-latency interactive apps, dev & test.
  - Volume Size: 1 TB – 16 TB.
  - Max IOPS/Volume: 16,000
  - Max Throughput/Volume: 250 MB/s
  - Price: \$0.10/GB-month



## Choose Region





13 / 30



14 / 30



15 / 30



16 / 30

# Configure Instance Details

Launch Instance wizard

1. Choose AMI 2. Choose Instance Type 3. Configure Instance Details 4. Add Storage 5. Add Tags 6. Configure Security Group 7. Review

### Step 3: Configure Instance Details

Configure the instance to suit your requirements. You can launch multiple instances from the same AMI, request Spot instances to take advantage of the lower pricing, assign an access management role to the instance, and more.

Number of instances: 1 [Launch into Auto Scaling Group](#)

**Purchasing option**

☐ Request Spot instances

**Network**  [Create new VPC](#)

**Subnet**  [Create new subnet](#)

**Auto-assign Public IP** ☐ Use subnet setting (Enable)

**Placement group** ☐ Add instance to placement group

**Capacity Reservation**

**Domain join directory**  [Create new directory](#)

**IAM role**  [Create new IAM role](#)

**Shutdown behavior**

**Stop - Hibernate behavior** ☐ Enable hibernation as an additional stop behavior

**Enable termination protection** ☐ Protect against accidental termination

**Monitoring** ☐ Enable CloudWatch-detailed monitoring  
Additional charges apply

**Tenancy**  [Additional charges will apply for dedicated tenancy](#)

[Cancel](#) [Previous](#) [Review and Launch](#) [Next: Add Storage](#)



17 / 30

# Configure Instance Details

Launch Instance wizard

1. Choose AMI 2. Choose Instance Type 3. Configure Instance Details 4. Add Storage 5. Add Tags 6. Configure Security Group 7. Review

### Step 3: Configure Instance Details

**Shutdown behavior**

**Stop - Hibernate behavior** ☐ Enable hibernation as an additional stop behavior

**Enable termination protection** ☐ Protect against accidental termination

**Monitoring** ☐ Enable CloudWatch-detailed monitoring  
Additional charges apply

**Tenancy**  [Additional charges will apply for dedicated tenancy](#)

**Elastic Inference** ☐ Add an Elastic Inference accelerator  
Additional charges apply

**Credit specification** ☐ Unlimited  
Additional charges may apply

**File systems** [Add file system](#) [Create new file system](#)

**Advanced Details**

**Metadata accessible**

**Metadata version**

**Metadata token response hop limit**

**User data** ☒ As text ☐ As file ☐ Input is already base64 encoded (Optional)

[Cancel](#) [Previous](#) [Review and Launch](#) [Next: Add Storage](#)



18 / 30

# Add Storage

Launch Instance wizard

1. Choose AMI 2. Choose Instance Type 3. Configure Instance Details 4. Add Storage 5. Add Tags 6. Configure Security Group 7. Review

### Step 4: Add Storage

Your instance will be launched with the following storage device settings. You can attach additional EBS volumes and instance store volumes to your instance, or edit the settings of the root volume. You can also attach additional EBS volumes after launching an instance, but not instance store volumes. [Learn more about storage options in Amazon EC2.](#)

| Volume Type | Device    | Snapshot              | Size (GiB) | Volume Type               | IOPS       | Throughput (MB/s) | Delete on Termination               | Encryption    |
|-------------|-----------|-----------------------|------------|---------------------------|------------|-------------------|-------------------------------------|---------------|
| Root        | /dev/xbd1 | snp-0049v0b3scv0yulc3 | 8          | General Purpose SSD (gp2) | 100 / 3000 | N/A               | <input checked="" type="checkbox"/> | Not Encrypted |

[Add New Volume](#)

Free tier eligible customers can get up to 30 GB of EBS General Purpose (SSD) or Magnetic storage. [Learn more about free usage tier eligibility and usage restrictions.](#)

[Cancel](#) [Previous](#) [Review and Launch](#) [Next: Add Tags](#)



19 / 30

# Add Tags

Launch Instance wizard

1. Choose AMI 2. Choose Instance Type 3. Configure Instance Details 4. Add Storage 5. Add Tags 6. Configure Security Group 7. Review

### Step 5: Add Tags

A tag consists of a case-sensitive key-value pair. For example, you could define a tag with key = Name and value = Webserver. A copy of a tag can be applied to volumes, instances or both. Tags will be applied to all instances and volumes. [Learn more about tagging your Amazon EC2 resources.](#)

| Key (128 characters maximum)                                   | Value (256 characters maximum) | Instances | Volumes |
|--|--------------------------------|-----------|---------|
| This resource currently has no tags                            |                                |           |         |
| Choose the Add tag button or click to add a Name tag.          |                                |           |         |
| Make sure your IAM policy includes permissions to create tags. |                                |           |         |

[Add Tag](#) (Up to 50 tags maximum)

[Cancel](#) [Previous](#) [Review and Launch](#) [Next: Configure Security Group](#)



20 / 30

# Configure Security Group

**Step 6: Configure Security Group**

A security group is a set of firewall rules that control the traffic to your instance. On this page, you can add rules to allow specific traffic to reach your instance. For example, if you want to set up a web server and allow internet traffic to reach your instance, add rules that allow unrestricted access to the HTTP and HTTPS ports. You can create a new security group or select from an existing one below. [Learn more](#) about Amazon EC2 security groups.

**Assign a security group:** Create a new security group

Select an existing security group

Security group name: launch-wizard-2

Description: launch-wizard-2 created 2020-10-27T11:38:56+02:00

| Type | Protocol | Port Range | Source | Description                |
|------|----------|------------|--------|----------------------------|
| SSH  | TCP      | 22         | Custom | e.g. SSH for Admin Desktop |

**Warning**

Rules with source of 0.0.0.0/0 allow all IP addresses to access your instance. We recommend setting security group rules to allow access from known IP addresses only.

[Cancel](#) [Previous](#) [Review and Launch](#)



21 / 30

# Launch Instance

**Step 7: Review Instance Launch**

Please review your instance launch details. You can go back to edit changes for each section. Click **Launch** to assign a key pair to your instance and complete the launch process.

**Warning**

Improve your instances' security. Your security group, launch-wizard-2, is open to the world. Your instances may be accessible from any IP address. We recommend that you update your security group rules to allow access from known IP addresses only. You can also open additional ports in your security group to facilitate access to the application or service you're running, e.g., HTTP (80) for web servers. [Edit security groups](#)

**AMI Details**

Ubuntu Server 20.04 LTS (HVM), SSD Volume Type - ami-0668a495a537da8b

**Instance Type**

| Instance Type | ECUs     | vCPUs | Memory (GiB) | Instance Storage (GiB) | EBS-Optimized Available | Network Performance |
|---------------|----------|-------|--------------|------------------------|-------------------------|---------------------|
| t2.micro      | Variable | 1     | 1            | EBS only               | -                       | Low to Moderate     |

**Security Groups**

| Security group name | Description                                       |
|---------------------|---|
| launch-wizard-2     | launch-wizard-2 created 2020-10-27T11:38:56+02:00 |

| Type | Protocol | Port Range | Source    | Description |
|------|----------|------------|-----------|-------------|
| SSH  | TCP      | 22         | 0.0.0.0/0 |             |

**Instance Details**

[Cancel](#) [Previous](#) [Launch](#)



22 / 30

# Create Key pair

**Step 7: Review Instance Launch**

Please review your instance launch details. You can go back to edit changes for each section. Click **Launch** to assign a key pair to your instance and complete the launch process.

**Warning**

Improve your instances' security. Your security group, launch-wizard-2, is open to the world. Your instances may be accessible from any IP address. We recommend that you update your security group rules to allow access from known IP addresses only. You can also open additional ports in your security group to facilitate access to the application or service you're running, e.g., HTTP (80) for web servers. [Edit security groups](#)

**AMI Details**

Ubuntu Server 20.04 LTS (HVM), SSD Volume Type - ami-0668a495a537da8b

**Instance Type**

| Instance Type | ECUs     | vCPUs | Memory (GiB) |
|---------------|----------|-------|--------------|
| t2.micro      | Variable | 1     | 1            |

**Security Groups**

| Security group name | Description                                       |
|---------------------|---|
| launch-wizard-2     | launch-wizard-2 created 2020-10-27T11:38:56+02:00 |

| Type | Protocol | Port Range | Source    | Description |
|------|----------|------------|-----------|-------------|
| SSH  | TCP      | 22         | 0.0.0.0/0 |             |

**Instance Details**

[Cancel](#) [Previous](#) [Launch](#)

**Select an existing key pair or create a new key pair**

A key pair consists of a **public key** that AWS stores, and a **private key file** that you store. Together, they allow you to connect to your instance securely. For Windows AMIs, the private key file is required to obtain the password used to log into your instance. For Linux AMIs, the private key file allows you to securely SSH into your instance.

Note: The selected key pair will be added to the set of keys authorized for this instance. Learn more about removing existing key pairs from a public AMI.

Create a new key pair

Key pair name: aws-ssh-key

[Download Key Pair](#)

You have to download the **private key file** (a .pem file) before you can continue. Store it in a secure and accessible location. You will not be able to download the file again after it's created.

[Cancel](#) [Launch Instances](#)



23 / 30

# Instance Ready to Launch

**Launch Status**

Initiating Instance Launches

Please do not close your browser while this is loading.

Creating security group... Successful

Authorizing inbound rules... Successful

Initiating launches...

[Cancel](#) [Previous](#) [Launch](#)



24 / 30

# List of Instances

The screenshot shows the AWS Management Console 'Instances' page. A table lists instances with columns for Name, Instance ID, Instance state, Instance type, Status check, Alarm status, Availability zone, Public IPv4 DNS, and Public IPv4. The instance 'i-00ca786c-fb8ef6c' is in a 'Running' state. Below the table, the details for this instance are displayed, including its Instance ID, Name, and various configuration options like Security, Networking, and Storage.



25 / 30

# Overview of Instance

The screenshot shows the 'Instance summary' page for instance 'i-078d572f100b2f8a9'. It provides a detailed overview of the instance, including its state (Running), type (t2.micro), and various configuration options like Security, Networking, and Storage. The page also includes a section for 'AWS Compute Optimizer' recommendations.



26 / 30

# Connect to Instance

The screenshot shows the 'Connect to Instance' page in the AWS Management Console. It features the 'EC2 Instance Connect' tab, which displays the instance ID, Public IP address, and User name. A red circle highlights the 'Connect' button, indicating the next step in the process.



27 / 30

# Command Line Console

The screenshot shows the 'Connect to Instance' page in the AWS Management Console, specifically the 'Session Manager' tab. It displays the instance ID, Public IP address, and User name. A red circle highlights the 'Connect' button, indicating the next step in the process.



28 / 30

## Other Examples

- Start, Stop, Terminate instance.
- Change Instance Type.
- Add Storage Volumes.
- Configure Security Groups.



## Connecting to the Instance

