

Hands-on Cloud Computing Services

Lezione 2

Gabriele Russo Russo

University of Rome Tor Vergata, Italy

A.Y. 2025/26



Recap


- ▶ Amazon Web Services: regions, services, ...
- ▶ Elastic Compute Cloud ([EC2](#))
 - ▶ Instance, AMI, Security Group
 - ▶ SSH, public/private keys
- ▶ Example web app: [Photogallery](#)
- ▶ Custom AMI for our app

Run Commands at Launch: `cloud-init` and User Data

- ▶ Creating a custom AMI allowed us to create new EC2 instances without manually configuring the application every time
- ▶ Alternative approaches?
- ▶ Cloud providers allow you to run commands when instances are launched:
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/user-data.html>
- ▶ In AWS, you can use the `User Data` option to specify:
 - ▶ a Bash script
 - ▶ `cloud-init` directives (<https://cloudinit.readthedocs.io/en/latest/>)

- ▶ Create a new EC2 instance with the custom AMI

Next step



- ▶ Provision **logically isolated** sections of the AWS cloud
- ▶ Define virtual networks (IP ranges, subnets, gateways,...)
- ▶ May create a Virtual Private Network (VPN) connection between your own datacenter and your VPC (**hybrid cloud**)
- ▶ **No additional charges** for creating and using the VPC itself.
- ▶ So far, we have used the **default VPC**

Amazon VPC: main building blocks

- ▶ In each AZ, we can define one or more **subnets**
- ▶ **Routing Tables** attached to subnets
- ▶ **Internet Gateway**

VPC Configuration: the hard way

- ▶ Create a new **Virtual Private Cloud (VPC)**
- ▶ We associate a block of (private) IP addresses to the VPC
 - ▶ Subnets will be created within this block of addresses
 - ▶ We can pick, e.g., 10.0.0.0/16
- ▶ We can create **subnets**: each subnet is associated with an Availability Zone (AZ)
- ▶ Let's pick an AZ and create a subnet (e.g., 10.0.1.0/24)
- ▶ If you want (for debugging), you can require that EC2 instances in the subnet are also assigned a public IP address
- ▶ Create an **Internet Gateway (IG)** to allow instances in the VPC to reach Internet; **associate** it with the VPC
- ▶ Create a **Route Table** for the VPC and **attach** it to the subnet(s)
- ▶ Add a new rule in the table: 0.0.0.0\0 – target: IG
- ▶ Repeat the above steps for each subnet you want.

VPC Configuration: the easy way

- ▶ AWS released a new UI to ease VPC configuration
- ▶ Most the elements you need automatically created along with the VPC
 - ▶ Subnets
 - ▶ Routing Tables
 - ▶ Internet Gateway (for public subnets)

- ▶ Create a new EC2 instance in one of the newly created public subnets
- ▶ Start with the custom AMI
- ▶ Make sure to enable the assignment of a public IP address

Elastic Load Balancing (ELB)

- ▶ ELB automatically distributes incoming traffic across multiple targets (e.g., EC2 instances, containers, and IP addresses) in one or more Availability Zones
- ▶ It monitors the health of its registered targets and routes traffic only to the healthy targets
- ▶ 4 types of ELB:
 - ▶ Application Load Balancer (layer 5)
 - ▶ Network Load Balancer (layer 4)
 - ▶ Gateway Load Balancer (layer 3)
 - ▶ Classic Load Balancer (legacy)
- ▶ We'll use the Application LB today

ELB Configuration

- ▶ Create an ELB instance listening for HTTP requests on port 80
- ▶ ELB needs a security group: configure one to accept traffic on port 80
- ▶ We must also create a **target group**, to which ELB forwards requests
 - ▶ Health check: use HTTP requests on port 80 with path /
- ▶ Create a few EC2 instances using our custom AMI in our subnets
- ▶ Register the instances with the target group
- ▶ Wait a few minutes (DNS...) and then try to connect at the ELB URL with the browser

Note:


- ▶ EC2 instances don't need a public IP address any more
- ▶ EC2 instances can now use a stricter security group:
 - ▶ Allowed source: 0.0.0.0/0 → <ID of ELB sec group>

ELB: Advanced Rules

- ▶ An ELB can have multiple **rules** associated to distribute requests
- ▶ Each rule can have one or more matching **conditions**
- ▶ e.g., you may use different rules for different types of HTTP requests

Auto scaling

- ▶ We want to dynamically provision the number of active instances
- ▶ Let's use the Auto Scaling service of EC2



Auto Scaling + Photogallery

- ▶ Before starting, terminate manually launched instances
- ▶ Create a Launch Template for Photogallery
- ▶ Create an Auto Scaling Group that uses the new Launch Template
- ▶ Specify the VPC and the subnets where new instances should be launched
- ▶ Enable load balancing, associating the group with our ELB
- ▶ Set minimum and maximum number of instances (e.g., 2 and 5)
- ▶ Set an auto scaling policy
- ▶ Verify that new instances are automatically created

Recap

- ▶ We have seen how to deploy a web app using:
 - ▶ EC2
 - ▶ ELB
 - ▶ Auto Scaling Groups
- ▶ **Problem:** infrastructure completely configured by hand
 - ▶ error-prone and difficult to reproduce

- ▶ [Command Line Interface](#) to interact with AWS
- ▶ Faster interaction compared to web console
 - ▶ e.g., EC2 instance created with a single command
- ▶ Installation: check the official docs for Linux/Win/macOS
 - ▶ <https://aws.amazon.com/cli/>
- ▶ [AWS CloudShell](#) provides an in-browser console where CLI commands are available (useful for quick commands)
- ▶ Alternatively, Windows users may prefer the [AWS Tools for PowerShell](#)
 - ▶ <https://aws.amazon.com/it/powershell/>

AWS CLI: Configuration

- ▶ AWS CLI can be configured by:
 - ▶ running `aws configure`, or
 - ▶ editing `~/.aws/config` and `~/.aws/credentials`
 - ▶ (slightly different paths on Windows)
- ▶ Key configuration options:
 - ▶ AWS Access Key ID and AWS Secret Access Key
 - ▶ default region to use (e.g., `us-east-1`)
 - ▶ output format (`json`, `text`)

AWS CLI: example (1)

Create a new security group in our VPC:

```
$ aws ec2 create-security-group \
  --group-name my-sg \
  --description "My security group" \
  --vpc-id <VPC_ID>
```

Set inbound traffic rules, e.g.:

```
$ aws ec2 authorize-security-group-ingress \
  --group-id <ID> \
  --protocol tcp --port 22 --cidr 0.0.0.0/0
```

We can see the properties of any SG:

```
$ aws ec2 describe-security-groups --group-ids <groupId>
```

AWS CLI: example (2)

Create an EC2 instance:

```
$ aws ec2 run-instances \
  --image-id <ID AMI> \
  --count 1 \
  --instance-type t2.nano \
  --key-name <MyKeyPair> \
  --security-group-ids <sgId> \
  --subnet-id <subnetId> \
  --associate-public-ip-address
```

We can associate the instance with a tag:

```
$ aws ec2 create-tags --resources <instID> \
  --tags Key=Name,Value=SDCC
```

AWS CLI: example (3)

We can get information about active instances:

```
$ aws ec2 describe-instances \
    --filters "Name=tag:Name,Values=SDCC"
$ aws ec2 describe-instances \
    --filters "Name=instance-type,Values=t2.nano"
```

To terminate the instance:

```
$ aws ec2 terminate-instances --instance-ids <ID>
```

IT Automation using Ansible

- ▶ Ansible delivers simple IT automation that ends repetitive tasks and frees up teams for more strategic work.
- ▶ Available on Linux and macOS: https://docs.ansible.com/ansible/latest/installation_guide/intro_installation.html
 - ▶ Windows users might use a Linux-based VM
- ▶ Agentless: no need to pre-install software on the target machines
- ▶ Define WHAT you want to achieve, instead of HOW
 - ▶ e.g., “Apache web server installed and started”
- ▶ Similar alternatives: Chef, Puppet, a bunch of Bash scripts, ...

Ansible: Key Concepts

- ▶ **Playbooks** (e.g., “deploy Photogallery”)
- ▶ **Tasks** (e.g., (“install Flask”))
- ▶ **Modules** (used to define single sub-tasks: e.g., file, archive, apt)
 - ▶ Built-in modules
 - ▶ Custom modules
- ▶ **Inventory** = hosts to be managed
 - ▶ Static
 - ▶ Dynamic

A playbook for Photogallery: inventory

- ▶ Create the inventory file `hosts.ini`
 - ▶ (You may even put `localhost` in the inventory for testing)
- ▶ One line per host
- ▶ Possibly organized into groups (e.g., `web`, `db`, ...)
- ▶ We can add params for SSH authentication on the same line

Inventory file

```
[web]
18.185.19.141 ansible_user='ec2-user' \
    ansible_ssh_private_key_file='/path/to/keypair.pem'
```

Simple test using the ping module:

```
$ ansible -i hosts.ini -m ping all
```

A playbook for Photogallery

To deploy Photogallery we need to:

- ▶ Upload application files (module: **copy**)
- ▶ Install dependencies (modules: **yum**, **pip**)
- ▶ Install systemd unit file to start server at boot (module: **copy**)
- ▶ Enable systemd service (module: **systemd**)

Check `deploy_gallery.yaml`

```
$ ansible-playbook -v -i hosts.ini deploy_gallery.yaml
# What happens if we try again?
$ ansible-playbook -v -i hosts.ini deploy_gallery.yaml
```

Ansible: Dynamic Inventory

- ▶ Ansible requires an inventory
- ▶ Not necessarily a static file
- ▶ AWS Inventory Source: run your playbooks using (a subset of) your EC2 instances as target hosts (e.g., filtered by tag)
- ▶ Requires Ansible 2.9+
- ▶ A plugin required, easy to install:

```
$ ansible-galaxy collection install amazon.aws
```

Ansible: AWS Dynamic Inventory

- ▶ Create a YAML file (name MUST end with `aws_ec2.(yml|yaml)`)
→ `galleryInventory.aws_ec2.yaml`

Test

```
ansible-inventory -i galleryInventory.aws_ec2.yaml --graph
```

Running the playbook

```
ansible-playbook -i galleryInventory.aws_ec2.yaml  
--private-key=path/to/key.pem -u ec2-user deploy_gallery.yaml
```

Ansible: More Advanced Stuff

- ▶ Groups and Roles
- ▶ Templates
- ▶ Ansible Tower / AWX¹
 - ▶ Share playbooks / delegate
 - ▶ Schedule workflows
 - ▶ Dashboards

¹<https://github.com/ansible/awx>

Amazon S3

AWS Storage Services

AWS offers various **storage** services, including:

- ▶ S3: Simple Storage Service
- ▶ EBS: Elastic Block Storage
- ▶ EFS: Elastic File System

Amazon S3

- ▶ Amazon Simple Storage Service (S3)
- ▶ Scalable object storage service
- ▶ Pricing: <https://aws.amazon.com/it/s3/pricing/>
- ▶ **Buckets and objects**

S3 for Photogallery

- ▶ Let's create a bucket using S3 console
- ▶ Bucket name must be unique across all AWS regions and accounts
- ▶ We can choose who can access objects and buckets: https://docs.aws.amazon.com/it_it/AmazonS3/latest/dev/example-bucket-policies.html
- ▶ For Photogallery, we want everyone to read objects

We can reference an object like this:

`https://BUCKETNAME.s3.amazonaws.com/FILENAME`

Using S3 through the CLI

```
$ aws s3 ls
$ aws s3 ls s3://mybucket
$ aws s3 cp prova.txt s3://mybucket/
$ aws s3 ls s3://mybucket
$ aws s3 rm s3://mybucket/prova.txt
```

Third-party clients also available: e.g., s3cmd

Hosting Static Web Content on S3

- ▶ Objects in a public bucket can be accessed through HTTP
- ▶ You can use S3 to host static web content
 - ▶ a static website
 - ▶ the frontend of a web application
- ▶ To enable web hosting on a bucket: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/EnableWebsiteHosting.html>