EC2 Storage types

Instance / Ephemeral Storage	Attached to the physical host running an instance
Elastic Block Store / EBS	Attached over the network
	preferable to instance storage in nearly all usage scenarios
	One can create a snapshot from an EBS.
	Once you have created a snapshot, you can then create additional EBS volumes that will be identical copies of the source snapshot. You could, for example, create a snapshot containing your database backups
	EBS volumes can persist after the instance has been terminated
EBS +ve:	EBS volumes is clearly preferable except in a few cases, such as when you need fast temporary storage for data that can be safely discarded.

Multiple volumes (of either type) can be attached to an instance

AWS EC2 Instance Types

Different combinations of CPU, memory, network bandwidth and even custom hardware differentiate AWS instance types.

General Purpose Instance	Balances computing, memory, and networking resources.
Compute Optimized Instances	Good for high-performance application servers, gaming servers, and web applications.
Memory Optimized Instances	To quickly deliver large dataset workloads
Accelerated Computing Instances	Boost data processing for graphics applications and streaming.
Storage Optimized Instances	For large datasets on local storage:
	- Large file systems
	- Data warehouses
	- Online transaction systems

AWS EC2 Instance Types (cont)

Selecting the right instance type:	drive the CPU to 100% using your application's load generator of choice. Now examine memory use: if you observe the instance running out of mem- ory before the CPU is at full throttle, switch to a higher-memory instance type. Con- tinue this process until you achieve a reasonable balance.

Amazon EC2 pricing

On-Demand	Short-term, irregular workloads that cannot be interrupted
Savings Plans	Reduce your compute costs by committing to a consistent amount of compute usage for a 1-year or 3-year term.
Reserved Instances	Billing discount applied to the use of On-Demand Instances in your account for a 1-year or 3-year term
Spot Instances	Are ideal for workloads with flexible start and end times
Dedicated Hosts	physical servers with Amazon EC2 instance capacity that is fully dedicated to your use.

By **datamansam**

cheatography.com/datamansam/

Published 28th April, 2023.
Last updated 26th April, 2023.

Purchasing EC2 Instances

On-demand:	Allocated by the hour and requiring no upfront commit-tment
Reserved:	Represent a pre-paid commit-tment on the part of a customer which is usually rewarded by AWS with very steep discounts, up to 75% of on-demand pricing.
Spot:	Requires no upfront committ-ment, and their pricing fluctuates according to the supply and demand of compute capacity.

Working with instances in the console

# Describe all of your own images in the US East region	aws ec2 describe--images --owners self --region us-east-1
# List the AMIs that have a specific set of key/value tags	aws ec2 describe--images --owners self --filters Name=tag:role,V-alues=webserver # List the AMIs that have a specific set of key/value tags Name=tag:enviro-nment,Values=produ-ction
# Basic invocation to create instances	aws ec2 run-instances--image-id ami-6d060707
# Another way to display information about your instance is with aws ec2 descri-beinstances, which will show much more detail	aws ec2 describe-ins-tance-status --instance-ids i-64a8a6fe --region us-east-1 --output text

AWS elastic compute workflow

1. Launch	a) Select a template with basic configs
	b) Specify security settings to control traffic in and out of instance
2. Connect	Users by loging in and accessing the computer desktop
3. Begin use	Run commands to:
	a) Install software
	b) Add storage
	c) Copy and organise files

EC2 Networking

launching an instance with the default networking configuration will give you an instance with a public IP address

Simple Config	Many applications will require nothing more complicated than enabling SSH or HTTP access
To more sophis itcated cofig	Amazon offers more-advanced solu- tions that can, for example, give you a secure VPN connection from your datacenter to a Virtual Private Cloud (VPC) within EC2.
VPC	A network dedicated to your account, isolated from other networks in AWS, and completely under your control.

EC2 Networking (cont)

You can create sub- nets and gateways, configure routing, select IP address ranges and define its security perimeter

Amazon makes a distinction between traffic destined for the public Internet and traffic that will remain on the internal EC2 network

Networking - Payload

Routing requires:

Address of sender

Payload or contents

Address of recipient

IP addresses:

Unique to each computer, binary

IPv4 notation - Usually binary IP are c

Introducing EC2

EC2:	Allows customers to rent computing resources by the hour in the form of virtual machines (known as instances) that run a wide range of operating systems.
Custom isation:	Instances can be customized by the user to run any software applications supported by their operating system of choice.

DB instances

Standard	
Memory Optimised	
Burstable performance	

Introducing Cloud Formation

A stack	A collection of AWS resources that you can manage as a single unit. I
	You can create, update, or delete a collection of resources by creating, updating, or deleting stacks.
Creating Stacks:	aws cloudformation create-stack --template-body file://example-stack.json \ --stack-name example-stack
Modifying Stacks:	$ aws cloudformation describe-stack-events--stack-name example-stack \ --output text
	$ aws cloudformation describe-stack-resources --stack-name example-stack \ --output text
Updating Stacks	. Update the stack.json file

Introducing Cloud Formation (cont)

	Update the running stack with aws cloudformation update-stack
	View with aws cloudformation describe--stack☐ event
Ensuring local copy of a stack matches the running version:	
- Get a JSON file with running template, cleaning output	aws cloudformation get-template
use diff to:	compare the local and remote versions

Cost Optimization

The ability to run systems to deliver business value at the lowest price point

Govern usage	Employing a checks-and-balance approach, you can innovate without overspending.
Monitor usage and cost	Establish policies and procedures to monitor and appropriately allocate
To decommission resources	Implement change control and resource management from project inception to end-of-life to u shut down or terminate unused resources to reduce waste

Evaluate cost when you select services:

Cost Optimization (cont)

Building Block AWS Services:	Amazon EC2, Amazon EBS, and Amazon S3
Application level:	Amazon RDS and Amazon DynamoDB
	Reduce or remove administrative tasks and operational overhead
Meet costs targets when selecting resources	Think type, Size and Number
Plan for data transfer charges:	A small yet effective architectural change can drastically reduce operational costs
Cost Explorer:	View and track your usage in detail
	Reserved Instance recommendations
Auto Scaling:	Match supply and demand

AWS Design Process

To identify any critical issues and areas that could be improved
Update answers as the architecture evolves
To facilitate meetings, provide:
• Print outs of any diagrams or design notes
• Action list of questions that require out-of-band research

By **datamansam**

cheatography.com/datamansam/

Published 28th April, 2023.
Last updated 26th April, 2023.

Identity and Access Management

Identity and Access Management (IAM) is the name given to the suite of features that let you manage who and what can access AWS APIs using your account.

The idea behind IAM is to separate users and groups from the actions they need to perform. You do this by creating an IAM policy, which is a JSON-formatted document describing which actions a user can perform.

Amazon Resource Names / ARM:	A globally unique identifier that references AWS objects
ARN Format:	arn:aws:service:region:account_ID:relative_ID
A permission:	a combination of two items: an action and one or more resources.
	Actions are namespaced strings that take the form service_name:Permission. All EC2-related permissions are prefixed with ec2:, such as ec2:DeleteSnapshot.

Create a set of access credentials that are authorized to perform only the specific actions required by the script:

Eg, an AMI policy with enough permissions to run the script that cleans old images and snapshots	Via four Boto function calls:

Identity and Access Management (cont)

One can use the command-line tools to create a user and attach a new policy to it, using the iam-usercreate and iam☐ useraddpolicy commands

Create a new user for this role, named ami-cleaner:	mike@ip-10-32-34-116:/tmp$ iam-usercreate -u ami-cleaner -k AKIAINAK6ZGJNUAWVACA cjJvjs79Xj/kvrJkLFpRrMAIMAx-EdQrJLGIQQD28
	The -k option says that we want to create a new access key and secret for this use. These get printed to stdout. The first item is the access key ID, and the second is the secret access key. Store these somewhere safe, as we will need them later.
Create an AMI policy and attach it to the user:	mike@ip-10-32-34-116:/tmp$ iam-useraddpolicy -u ami-cleaner -p ami-cleaner -e Allow -r "*" -a ec2:DescribeImages -a ec2:DeleteSnapshot -a ec2:DeregisterImage

Identity and Access Management (cont)

The -e and -r arguments state that we are creating an Allow policy that applies to all resources (the asterisk after the -r option). Finally, we specify a list of actions that will be allowed, each preceded by an -a flag. You can specify as many permissions as you need.

Referencing resources in IAM policies

The Resource attribute of an IAM policy lets you control exactly which resources an action can be performed on. In the previous example, the policy granted the user permissions to delete any EBS snapshot owned by this account. What if you want a more granular policy that applies only to a subset of resources?

ARNs are used to globally identify AWS resources. Used in IAM policies, they let you control exactly which resources are included when grant- ing or denying permissions.

An IAM policy that allows users to perform any action on S3 buckets, with the exception of the one containing your backups. We do this by creating a policy containing two statements. The first grants the user all S3-related permissions, allowing them to be performed on any resource. The second statement denies all S3-related permissions, but only on the protected buckets

Identity and Access Management (cont)

{ "Statement": [{ "Action": ["s3:*"], "Effect": "Allow", "Resource": ["*"] }, { "Action": ["s3:*"], "Effect": "Deny", "Resource": ["arn:aws:s3:::db-backups"] }] }

Next, create the policy using the command-line tools or Management Console. If using the Management Console, you can create the policy as follows: 1. Navigate to IAM Users. 2. Select an existing User or Group. 3. Click Attach User/Group Policy. 4. Select Custom Policy. 5. Paste the text into the Policy Document box

Dynamic Policies:

Conditions can be used to create dynamic IAM policies that behave differently, depending on one or more factors. The attributes of the request (such as the ARN of the requesting user or the source IP address) can be used in Boolean expressions to control whether a request should be allowed or denied.

Attributes on which you can base your conditions are as follows:	• Time of day • Source IP address • Whether the request is being made using HTTP or HTTPS

Limitations of IAM policies

Some AWS resources do not use ARNs, and can therefore not be explicitly managed by IAM policies.	Because EC2 instances do not have ARNs, there is no way to reference a specific EC2 instance from an IAM policy

Identity and Access Management (cont)

Whenever you refer to EC2 permissions in a policy, the resource will be *, which means it will apply to every instance owned by your AWS account.

IAM Users and Groups

The user can be assigned one or more IAM policies, which specify the actions the user is allowed to perform.

Users can be placed in groups. When an IAM policy is assigned to a group, all members of that group inherit the permissions designated by the IAM policy

IAM is a global AWS service, meaning it is not tied to any particular region. An IAM user will be able to access APIs in any region

Map AWS groups to specific roles within your organization, and apply the policy to the group instead

Amazon’s CloudTrail service keeps track of the API calls made by
users in your account. You can use this to review the full history of
AWS API calls that have been made by your account, whether they
came from the Management Console, cli tools, or services like
CloudFormation

By **datamansam**

cheatography.com/datamansam/

Published 28th April, 2023.
Last updated 26th April, 2023.