

IAM

- Users, Groups (users heritage its policies), Roles (for aws resources) and Policies (json)
- Global
- Root Account (never use it) + MFA
- By default new users have no permissions
 - o Programmatic access
 - o AWS Management console access
- IAM password policy
- Roles:
 - o AWS Service Role - The usual one, and the one we are interested in
 - o AWS service-linked role: for Alexa
 - o Role for cross-account access: allow IAM users to access to another AWS accounts
 - o Role for identity provider access: grant access from Cognito, or OpenID (facebook, google, amazon), SAML, etc

S3

- Key-Value Object Storage. Files from 0 byte to 5TB. Unlimited storage
- S3 buckets: universal namespace. Default: max 100 buckets/account
- Read after write for PUTs of new objects
- Eventual consistency for overwrite PUTS and DELETES

S3 object consists on:

- Key: nombre del fichero
- Value: el contenido del fichero (secuencia de bytes)
- Version ID
- Metadata
- Subresources
- Access Control List

S3 tiers

- **S3 standard:** Objeto mínimo 0 bytes
- **S3 IA:** accedes 1 vez al mes (o cada 6 meses). Pero necesitas acceso rápido
 - o Objeto mínimo 128Kb. Es la opción más barata de S3
- **RRS:** para ficheros que puedes permitirte perder. Thumbnails
- **Glacier:** para archivar. Tardas de 3 a 5h en recurrarar un fichero
 - o Restauras via S3 API o via la consola de AWS.

	Standard	Standard - Infrequent Access	Reduced Redundancy Storage
Durability	99.999999999%	99.999999999%	99.99%
Availability	99.99%	99.9%	99.99%
Concurrent facility fault tolerance	2	2	1

Bucket URL formats:

<http://s3-<region>.amazonaws.com/<bucket>> o <http://<bucket>.s3-<region>.amazonaws.com>

Versioning

- No puede deshabilitarse, solo suspenderse
- Cada update es un fichero por sí mismo, con su propio ID
- Eliminar un fichero es marcarlo (**delete marker**) como eliminado > desaparece del bucket, no del histórico > [Sólo el propietario del bucket puede eliminarlos de verdad](#)
- Puedes habilitar MFA para los deletes

Cross region replication (CRR)

- Require **“versioning”** habilitado.
 - o Permite subconjuntos via prefijos. [Tb replica metadatos y ACL](#)
- Al subir algo nuevo (o update) al bucket, se replica a otro bucket (en otra región) – tb requiere **versioning** pero acepta otro tipo de S3 (IA, RRS...). Requiere IAM roles.

Lifecycle & Glacier

- Sin versioning
 - o **30 días de S3 > IA** (sólo para objetos mayores de 128KB)
 - o **30 días de IA > Glacier**
- Con versioning
 - o Tengo 2 LC, uno para el objeto actual y otro para las versiones antiguas

S3 Security & Encryption

- Por defecto los buckets son privados
- Control de acceso via bucket policies (aplica a todos los objetos) o ACL
- Puedes habilitar logging > lo guarda en otro bucket
- Encryption
 - o In transit (SSL/HTTPS) – [SSL/HTTP endpoints using HTTPS protocol](#)
 - o At rest
 - Server Side Encryption (SSE)
 - S3 Managed Keys (SSE-S3). Amazon se encarga de todo.
 - AWS Key Management Service (SSE-KMS) **Permite Audit Trail**
 - Customer Provided Keys (SSE-C): Tú controlas las claves
 - Client Side Encryption

S3 Transfer Acceleration

- Usa las edge locations de CloudFront para subir los datos desde el más cercano a ti
 - o Coste adicional. Debes usar la URL proporcionada para esas transferencias

S3 Static Website Hosting

- Si usas Route53 con S3, el nombre del bucket debe ser el del dominio (sin el “.com”)
- **[http://<bucketname>.s3-website-<region>.amazonaws.com](#)**
- Puedes especificar index/error pages y redirect rules

CloudFront

- Edge location: caché, TTL (default 24h), puedes habilitar la escritura/update en edge locations que updatean el origen
 - o Puedes elegir "Allowed HTTP methods" (GET, HEAD, PUT, DELETE...)
- Origin (permite múltiples orígenes para la misma distribución)
 - o S3 bucket: puedes restringir el bucket para que sólo se pueda acceder desde el CDN -> Origin Access Identity
 - o EC2 instance
 - o ELB
 - o Route53
 - o Fuera de AWS
- Distribution
 - o Web distribution: para websites
 - o RTMP: media streaming

S3 multipart upload API > [abort or failed uploads via lifecycle policies](#). Puede usarse con tx acc

- [Recommended for files > 100MB](#)

Storage Gateway

- VM que instalas en tu datacenter y replica a S3.
- 3 tipos
 - o **Gateway Storage Volumes:** tus datos en local, SGW replica a S3 (bkp)
 - o **Gateway Cached Volumes:** tus datos en S3, SGW sirve de caché local
 - o **Gateway Virtual Tape Library (VTL):** reemplaza los bkps en cinta > usa S3

Import/Export

- Actualmente reemplazado por Snowball. Permite:
- Exportar desde S3
- Importar a S3, Glacier y EBS

Snowball

- Importar/Exportar hacia/desde S3.
- Snowball: Petabyte scale data transport solution
- Snowball edge: + compute capabilities. i.e gather data during a flight
- Snowmobile: el camión. Exabyte scale

EC2

Pricing

- On demand
- Reserved: 1 or 3 years. Predictable usage or Reserved Capacity
- Spot: flexible start/end, only feasible at low prices, urgent compute needs
 - o Si la termina AWS, no pagas por esa fracción de hora
- Dedicated hosts: Por hora o Reserved. Licencias o for Regulatory Requirements

Types

- Dr Mc Gift Px

EBS (Elastic Block Storage)

- General Purpose SSD (gp2). 3iops/GiB max 10K iops
- Provisioned iops SSD (io1). Por si necesitas más de 10K iops (hasta 20K)
- Throughput optimized HDD (ST1). Frequent Access. Large amount of data in sequence as Data warehousing, log processing. Cannot boot
- Cold HDD (SC1). Less frequent access. Typical: fileserver. Cannot boot
- Magnetic (Standard). Infrequent access, lowest cost
- Por defecto: root volumen terminated al terminar la instancia
- Los volúmenes deben estar en la AZ de la instancia que los quiere usar
 - o EBS guarda copias redundantes dentro de la misma AZ

EBS: upgrading volumes (cambiar tamaño o tipo)

- BEST PRACTICE: parar instancia, detach, hacer snapshot, crear new volumen, attach.
- EBS pueden updatearse on the fly (excepto magnetic standard)
 - o Sólo un cambio en 6 horas
- El tamaño sólo puede incrementarse (incluso desde snapshot)

RAID & EBS

- Aumentar iops = Raid 0 (stripped) o 10
- **Application Consistent Snapshots:**
 - o Necesita 1) parar escrituras a disco desde la aplicación 2) flush caché
 - o 3 métodos para hacer esto:
 - Freeze the filesystem
 - Unmount the RAID array
 - (BEST OPTION) Parar la instancia, tomar snapshot, iniciar instancia

EBS Snapshots:

- Puedo: Crear Volumen, AMI, copiarlo a otra región y/o crear una copia "cifrada"
- No puedo eliminar un snapshot usado por una AMI (creada a partir de él)
- Los snapshots se almacenan en S3, y son incrementales ([allow point-in-time recover](#))

Encrypt Root device volume and create AMI

- No puedo crear un snapshot cifrado de un volumen no cifrado
- Los snapshots hechos de volúmenes cifrados, están cifrados automáticamente
- Los volúmenes restaurados desde snapshot cifrados, están cifrados automáticamente
- Sólo puedes compartir AMIs NO cifradas (con otras cuentas AWS o públicamente)
- Las AMIs son "por región" pero puedo copiarlas

EBS root vs instance (ephemeral) storage

- Si el root device es EBS, éste creó lanza desde una AMI creada de un snapshot EBS
- Si es instance store, éste se creó desde una AMI creada desde un template en S3 (slow)
- Las instancias con instance storage no se pueden parar (si el host falla, la info se pierde)
- Puedes escoger no terminar los EBS root volumes, pero NO los instance storage.
- [No puedes desatachar el root EBS sin parar la instancia, claro](#)

Security Groups

- Por defecto: inbound denied, outbound allowed
- Cambios aplicados inmediatamente
- Son stateful: crean reglas (no visible) para el tráfico relacionado

ALB/ELB y Healthchecks -> self-sanitization of instances

- Tienen su propio security group
- LB asociado a una VPC. Puede (debe) trabajar en varias AZ
- No tienen IP, sólo un DNS record
- [Cross-Zone enabled = Balancea entre instancias, independientemente de las AZ](#)
- ELB (capa 4)
 - o No permiten instancias creadas desde Amazon DevPay site
 - o SSL Termination: has de instalar el certificado en el ELB
 - o Puedes loggear la actividad con CloudTrail
- ALB (capa 7) + [Barato](#)
 - o Internet facing o internal
 - o Routing > target groups = path based routing! (ie. /a > target1, /b > target2)
 - o Healthcheck opcionalmente puede chequear el HTTP success code
 - o Parar SSL termination en las instancias

CloudWatch for EC2

- Default metrics on EC2 instances: CPU, disk, Network, Instance status
- Standard monitoring (5min) vs detailed (1min)
- Dashboards, alarms, events (responde a cambios en los recursos de AWS) and logs (requiere un agente instalado en la instancia. Permiten agregar y almacenar logs)
- Cloudwatch (monitoring y logging) VS CloudTrail (para auditar)
- Tipos de alarma: OK, Alarm, insufficient-data

Userdata & Metadata

- Bootstrap scripts: user data section (max 16KB)
- Instance Metadata: <http://169.254.169.254/latest/meta-data/>

Launch configuration & ASG

- Launch configuration: plantilla con la creación de imágenes
- ASG: size, VPC y subnets donde crear las instancias, ELB, Healthcheck (ELB o EC2)
 - o + Scaling Policies: min/max & increase/decrease when...
 - o termination: selects AZ with most instances > delete the one using the oldest lc
 - o cooldown: seconds after another scaling event can happen

EC2 termination protection deshabilitado por defecto

EC2 Placement groups

- Grupo lógico de instancias que necesitan **low latency** y/o **high network throughput**
 - o 10Gbps. Misma AZ
- El nombre del PG debe ser único en tu cuenta AWS
- Sólo para cierto tipo de instancias (cpu, ram, storage y gpu)
- No puedes juntar PG. Tampoco mover una instancia de un PG a otro.

EFS

- Soporta NFSv4 y miles de conexiones simultáneas
- Petabytes. Data stored in multiple AZ in a región
- **Read after write consistency**
- Tiene su propio sg para cada punto de montaje = subnet = AZ
- Puede almacenar datos de una bbdd (al igual que EBS)

Lambda

- Puedes usarlo:
 - o Event-drive compute service: en respuesta a eventos
 - o En respuesta a HTTP requests via API Gateway
- Lenguajes: Java, NodeJS, Python, C#
- Triggers:
 - o API Gateway
 - o IoT
 - o Alexa
 - o CloudFront
 - o CloudWatch
 - o CodeCommit
 - o Cognito
 - o DynamoDB
 - o Kinesis
 - o S3
 - o SNS
- Máxima duración 5 min
- Las ejecuciones son independientes
- Escala horizontalmente (scale out) automáticamente

API Gateway

- Publish, maintain, monitor and secure APIs to EC2 or Lambda
- You can enable API caching to cache (for a TTL) the API response
- You can throttle (estrangular) API GW to prevent attacks
- You can log results to CloudWatch
- CORS (Cross-Origin Resource Sharing) > permite servir contenido de un dominio diferente al original

Route53

- ELB do not have IPv4, you resolve to them via DNS name
- Understand Alias (you can resolve individual AWS resources) vs CNAME
- Routing policies:
 - Simple (default): round robin
 - Weighted: A/B
 - Latency: lowest network latency (ms) to a region > latency
 - Failover: active/passive setup -> healthchecks
 - Geolocation: latency & show a geo-customized web
- Default limit of 50 domain names (can be increase contacting support)

Databases

RDS for OLTP

- Have to select instance type, EBS size (max 6TB/16TB for SSD), VPC, etc.
 - o [SQL Express max 300GB disk size](#)
- **Backups**: Automated (enable 1 by default 1-35 days) VS Database snapshots > impact performance! -> Backup window ([changes to it applied immediately](#))
 - o [Automatic backups are deleted when terminate \(only latest snap could be\)](#)
- **Encryption** only at creation time!!! [Not even from snapshots \(I think\)](#)
- **Multi-AZ**: only for disaster recovery. Does not improve performance. AWS Handles failover -> Sync replication
- **Read Replica**: Read performance. Requires auto backups on. Max 5, same AZ. Async
 - o [Available for MySQL and PostgreSQL engines](#)
- [Permite aplicar particionado de tablas para usar varias instancias RDS](#)
- Aurora: 5x faster than MySQL
 - o Maintains 2 copies of physical data in 3 AZ (min 6 copies)
 - Can fail 2 for writes, 3 for reads
 - o 2 Types of replica: Aurora (max 15, fault tolerance) & MySQL Read Replicas

DynamoDB – NoSQL

- Really scalable (no downtimes!), fast (SSD) and flexible
- Spread across 3 data centers
- Eventual Consistent Reads (if you can wait 1 second)
- vs Strong Consistent Reads (if you can't) -> increases cost
- Very cheap for reads
- Provisioned capacity = ios per table
- [Exists an option for Cross Region Replication](#)

Redshift for OLAP (& BI)

- Single node (160G)
- vs Multi-node, consists on
 - o Leader Node
 - o up to 128 Compute Nodes
- Fast because
 - o Columnar data storage ([block size = 1MB](#))
 - o Advanced compression (by columns)
 - o Massive Parallel Processing (MPP) across all nodes

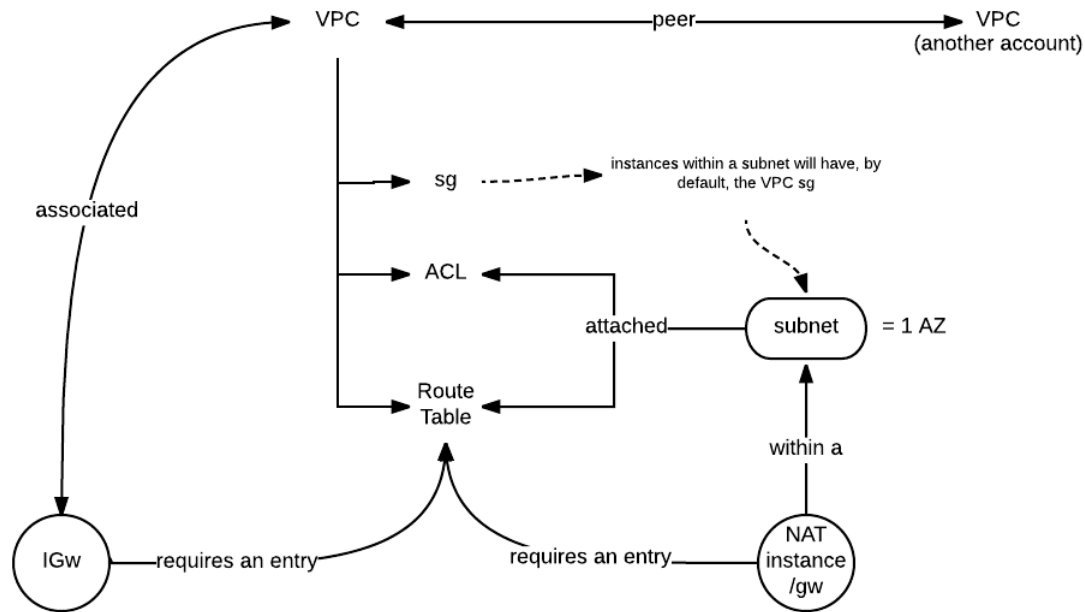
Elasticcache

- Memcached and Redis

[SSD better performance than magnetic for DBs in EC2 instances](#)

[RDS troubleshooting > look for "error nodes" in XML RDS API response](#)

VPC



VPC

- Private datacenter
- Max 1 IGW per VPC. After created, detached
 - o Route table has to have a route through IGW
- VPC peering, even with another AWS accounts (NO TRANSITIVE PEERING)
 - o IP ranges cannot overlap!!
- Custom VPC creates
 - o Default ACL > all denied by default
 - o Default SG
 - o Main Route Table > allow local (private) connections > **so by default, all subnets within the VPC can communicate to each other**
- By default, max 5 VPCs per region
- **Instances in default VPC will have public and private IP**
- **VPC endpoints to access to AWS resources**
- **VPC Flow logs: capture traffic within the VPC and sends it to CloudWatch**

Subnet

- 1 subnet = 1AZ
- Only can be attached to 1 ACL, and associated to 1 Route Table
- Public means the route table where is associated has an IGw, and its instances has a public IP

NAT

- o To allow instances within a private subnet to reach internet (for yum, i.e.)
- o Be placed in a **public subnet** (so with an IGw attached)
- o Needs an entry in the route table associated with the private subnet
- Nat instance is just a regular EC2 instance with a specific AMI
 - o Needs a public IP
 - o Needs disable "source/destination check"
 - o HA via ASG, multiple subnets and a script to automate failover
 - o Throughput depends on instance type
- Nat GW
 - o Scale automatically up to 10Gbps, across a single AZ

ACLs

Sg	ACL
Instance level (1 st)	Subnet level (2 nd)
Allow rules	Allow/Deny
Stateful	Stateless
All rules evaluated before deciding	FW: Rules in asc order > first match
Only applies to the instance if attached	Applies to all instances in the subnet

- Ephemeral ports for outbound connections (1024-65535)
- Your VPC automatically have a default ACL, with by default all inbound/outbound traffic is enabled
- But when you create your custom network ACL, all inbound/outbound traffic is denied

Application Services

SQS: pull. queue. message oriented API

- Simple Queue Service: Pull queue message system
- To **decouple** your components < EXAM!!
- Message size 256KB any format (text, json, xml)
- Messages in queue from 1min to 14 days. Default 4 days
- **Visibility timeout**: tiempo que tiene un consumer para procesar el mensaje (max 12h)
 - o Si da timeout, el mensaje vuelve a la cola > Puede duplicarse!
- **Long Polling**: en lugar de preguntar cada X seg si hay mensajes, preguntas y te avisa al entrar mensajes, o cuando de el long poll timeout ([ReceiveMessageWaitTimeSeconds>0](#))
- 2 tipos: default (**puede haber duplicados, no en orden**) y fifo

SWF: task oriented API

- Simple Workflow Service. Can include **human interaction**
- **Workflows max 1 year**
- **A task is assigned only once, never duplicated, and in order**
- SWF tracks all events. With SQS you have to implement your app-level tracking
- Parameters in JSON
- "Domains" are a collection of related workflows.
 - o Includes "workflow starters", "deciders" and "activity workers".

SNS: push. message oriented API

- Simple Notification Message: publish-subscribe service
- **mobile push notifications, Email/Email-JSON, SMS, SQS or Lambda**
- **SNS topics**: access points for clients to allow to subscribe to notifications (also HTTP(S))
- Data format in JSON

Elastic Transcoder: media converter

Kinesis

- Stream: consists on **shards. Data retained** max 7 days (default 1)
 - o Producer > Shards within the stream > Consumers
- Firehouse: no shards, streams or consumers. Data send to S3. **Optional Lambda** analysis
 - o Producer > Firehouse (optional Lambda) > S3
- Analytics: encima de Streams/Firehouse añade SQL analytics

Withpapers: Security

Shared security model

- AWS is responsible for the security config of its **managed services** products (DynamoDB, RDS, Redshift, EMR, WorkSpaces, etc.) and the underlying infra
- YOU: IAAS (EC2, VPC, S3) are under your control
- YOU are responsible for account & user access.
 - o Recommend MFA, SSL/TLS for communications and CloudTrail for user activity logging

Storage Decommissioning

- AWS uses DoD 5220.22 (National Industrial Security Media Sanitization) or NIST 800-88 (Guideless for Media Sanitization) to destroy data.
- Magnetic storage devices are physically destroyed

Network Security

- You can connect to AWS via HTTP or HTTPS using SSL
- VPC allows to use IPSec VPNs to tunnel between AWS and your datacenter
- AWS network is segregated from the Amazon Corporate (.com) network

Network Monitoring & Protection

- By default, AWS provides protection for
 - o DDoS
 - o Man in the middle
 - o IP Spoofing: the AWS host-based firewall will not allow instances to send traffic with a source IP or MAC other than its own.
 - o Port Scanning
 - o Packet Sniffing by other tenants (inquilinos)
- Unauthorized port/vulnerability scans by EC2 users are a violation of AWS Acceptable Use Policy. You may request permission before!

AWS Credentials

- Passwords
- MFA
- Access Key
- Key Pairs: SSH login to EC2. Cloudfront signed URLs
- X.509 Certificates: SSL certificates for HTTPS/ SOAP-based requests to AWS API

Trusted Advisor

- Inspects your AWS environment and makes recommendations to
 - o Save money
 - o Improve performance
 - o Close security gaps
 - o Fault Tolerance
- Provides alerts of common security misconfigurations

Instance Isolation

- Instances running on the same box, are isolated from each other via the Xen hypervisor.
 - o AWS firewall in the hypervisor layer between physical and EC2 NICs
- Physical RAM is separated using similar mechanisms
 - o Memory allocated to guest is scrubbed (set to zero) when unallocated.
- Instances have no raw access to disk, but a virtual disk.
 - o AWS automatically resets (disk zeroing) every customer's block of storage

Other considerations

- Gest OS:
 - o virtual instances are completely controlled by you. No backdoors for AWS!
 - o good security practice: EBS volumes and snapshots encrypted with AES-256
- ELB: Supports SSL Termination on the LB > instances can identify the source IP address
- Direct Connect: dedicated connection from your datacenter to your AWS VPC, using 802.1q VLAN standard, allowing you to connect to AWS public resources (S3) and private ones (EC2 in a private subnet)

Exam feedback

Virtualization types

- Paravirtual (PV)
- Hardware Virtual Machine (HVM)
 - o Better performance
 - o Can take advantage on hardware extensions and run in top of hw

AD

- Directory Service's AD Connector: let's you connect your existing AD to AWS
- Simple AD: inexpensive AD compatible with the common AD features
- You can authenticate with AD to AWS using SAML
 - o Authenticate to AD first, then to STS

AWS Organization & Consolidation Billing

- Account Management service to manage multiple AWS accounts from a central location
- Consolidated billing: 1 billing-only account. Up to 20 linked accounts. Global discounts

Resource Groups & Tagging

- Groups resources that share one or more tags

Security Token Services (STS)

- Federation (typically AD) – means join groups –
 - o Uses SAML
 - o Allows users to login to AWS without IAM credentials (but AD)
- Federation with Mobile Apps
 - o Uses Fb, Google, OpenID to login
- Cross Account Access

Workspaces

- VDIs. Are persistent.
- Runs Win7. By default users are local admins (allow to install applications)
- All data on D: is backed up every 12h

ECS

- ECR: EC2 Container Registry
- ECS Tasks definitions are JSON files describing one or more containers that conform your application (include CPU, RAM, links, etc)
- ECS service is like ASG using Task Definitions
- Clusters (region specific) are logical groups of container instances to place tasks in
- Service Scheduler: ensures a specific number of tasks is constantly running (ELB reg)
- Custom Scheduler: third party
- ECS Agent (docker agent)
- EC2 uses IAM roles to access ECS (Security groups still at host (EC2) level)
- ECS tasks uses IAM roles to access services and resources

More info:

https://acloud.guru/forums/aws-certified-solutions-architect-associate/discussion/-KSDNs4nfg5ikp6yBN9l/exam_feedback

Extra

- 44 AZ, 17 regions
 - o AZ names are assigned randomly per account!!!
- For new AWS accounts > max 20 EC2 instances per region
- 4 support levels: basic, developer, business, enterprise
- Por defecto, max 5 EIP por región > las EIP estarán atachadas a la instancia hasta que explícitamente las detaches (no se detachan si la instancia se para)
- CloudTrail permite registrar el histórico de llamadas a la API de AWS
- AWS Config permite guardar el histórico de cambios en las configuraciones de recursos de AWS > y enviar notificaciones de cambios via SNS
- 1GiB <= EBS size <= 16TiB
- RPO (Recovery Point Objective): datos que estoy dispuesto a perder (ej. 1h)
RTO (Recovery Time Objective): tiempo en volver a dar servicio (ej. 20j)

EXAM QUESTIONS

- for sure: 1 subnet = 1AZ -> you cannot spread 1 subnet across AZ
- NO VPC Transitive peering: if A – B – C, A cannot talk with C through A. You have to create a peer between A and C.