Day 3 of Preparation, 25th August, 2026.

Schedule

- 7am: Get Ready for Work - Shower, Breakfast, Hygiene Routines.
- 8am: Leave for Work.

Agile, Part 1

- Agile Manifesto
  - Individuals and Interactions
  - Working Software
  - Customer Collaboration
  - Responding to Change
- Agile Principles
  - Satisfy the Customer through Early and Continuous Delivery of Valuable Software
  - Welcome Changing Requirements, even late in development
  - Deliver Working Software frequently (weeks rather than months)
  - Close, Daily Cooperation between Business People and Developers
  - Projects are built around motivated individuals, who should be trusted
  - Face-to-face conversation is the most efficient and effective method of conveying information within a development team
  - Working software is the primary measure of progress
  - Sponsors, Developers, and Users should be able to maintain a constant pace indefinitely
  - Continuous attention to technical excellence and good design enhances agility
  - Simplicity - the art of maximizing the amount of work not done - is essential
  - The best architectures, requirements, and designs emerge from self-organizing teams
  - At regular intervals, the team reflects on how to become more effective, then tunes and adjusts its behavior accordingly
- Agile Methodologies
  - Scrum
  - Kanban
  - Extreme Programming (XP)
  - Lean
  - Crystal
  - Feature Driven Development (FDD)
  - Dynamic Systems Development Method (DSDM)
  - Rational Unified Process (RUP)
  - Agile Modeling
  - Agile Unified Process (AUP)
- Scrum
  - Product Owner
  - Scrum Master
  - Development Team
  - Product Backlog
  - Sprint Backlog
  - Sprint
  - Daily Scrum
  - Sprint Review
  - Sprint Retrospective
  - Burndown Chart
  - Velocity
  - Definition of Done
  - Acceptance Criteria
  - User Stories
  - Epics
  - Tasks
  - Spikes
  - Bugs
  - Technical Debt
  - Refactoring
  - Continuous Integration
  - Continuous Delivery
  - Continuous Deployment
  - Test-Driven Development (TDD)
  - Behavior-Driven Development (BDD)
  - Acceptance Test-Driven Development (ATDD)
  - Pair Programming
  - Code Review
  - Coding Standards
  - Version Control
  - Branching Strategy
  - Merge Conflicts
  - Release Management
  - Staging Environment
  - Production Environment
  - Monitoring
  - Logging
  - Alerting
  - Incident Management
  - Post-Mortem
  - Root Cause Analysis
  - Blameless Culture
  - Psychological Safety
  - Retrospective
  - Sprint Retrospective
  - Start/Stop/Continue
  - Mad/Sad/Glad
  - Sailboat
  - Timeline
  - Lean Coffee
  - Dot Voting
  - Action Items
  - Follow-Up
  - Continuous Improvement
  - Kaizen
  - Kanban Board
  - Work In Progress (WIP) Limits
  - Cycle Time
  - Lead Time
  - Throughput
  - Cumulative Flow Diagram
  - Control Chart
  - Pull System
  - Push System
  - Value Stream Mapping
  - Waste
  - Muda
  - Overproduction
  - Waiting
  - Transportation
  - Inventory
  - Motion
  - Defects
  - Overprocessing
  - Underutilized People
  - Poka-Yoke
  - Andon
  - Jidoka
  - Just-in-Time (JIT)
  - Heijunka
  - Takt Time
  - Kaizen Blitz
  - Gemba
  - Genchi Genbutsu
  - Hoshin Kanri
  - Policy Deployment
  - Catchball
  - Nemawashi
  - Ringi
  - Ringi-sho
  - Ringi-kessho
  - Ringi-tsumi
  - Ringi-mawashi
  - Ringi-ketsu
  - Ringi-ketsu-sho
  - Ringi-ketsu-sho-tsumi
  - Ringi-ketsu-sho-tsumi-mawashi
  - Ringi-ketsu-sho-tsumi-mawashi-ketsu
  - Ringi-ketsu-sho-tsumi-mawashi-ketsu-sho
  - Ringi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi
  - Ringi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi
  - Ringi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu
  - Ringi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho
  - Ringi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi
  - Ringi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi
  - Ringi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu
  - Ringi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho
  - Ringi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi
  - Ringi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi
  - Ringi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu
  - Ringi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho
  - Ringi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi
  - Ringi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi
  - Ringi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu
  - Ringi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho
  - Ringi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi
  - Ringi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi
  - Ringi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawashi-ketsu-sho-tsumi-mawash

Day 2. 28th August. VPC Focus.

Schedule:

- Warm-up
- V1: VPC Anatomy
- V2: Gateways and Network Balance Translation
- V3: Security group vs network access control list
- V4: Leaving the VPC - peering, VPC endpoints, Direct Connect vs VPN, Transit Gateway.

Warm-up

- An Amazon EBS volume is a block device or ONE instance in ONE AZ (Availability Zone).
  - every instance sees only its own frames.
  - EBS Multi-attach exists, but it's io1/io2 only. only AZ
- Amazon EFS is NFSv4, one filesystem, a mount target in each AZ, POSIX lock (mac).
- What are the 2 filesystem operations an object store cannot give that this filesystem workload needs? What would need to change for S3 to be the solution?
  - concurrent file modification IS what EFS is for. S3 and EBS **can't** do concurrent file modification.
- Recall the equation 10 x 10 = 1MRPS = days of transfer.

V0: VPC Slides.

- availability zones, route tables. Private subnet, Public subnet. Internet Gateway, Virtual Private Gateway.
- Region encapsulates VPC, which includes (maybe) encapsulates Route Table, which then encapsulates Private Subnet and Public Subnet.
  - Private Subnet must reside entirely within one AZ only.
    - one subnet cannot span to two or more AZs
  - Public Subnet can have multiple subnets in the same Availability Zone
- We can also bind a Network ACL on the Private Subnet and Public Subnet, this seems to increase security somehow.
  - Network ACL = ?
- In terms of services, **backend** systems like DATABASES or APPLICATION servers are not meant to be accessed publicly.
  - EFS, RDS, Pitu, Private EC2 Servers
  - hence on the **private subnet**
- We can have Public EC2 Web Servers (part of a security group), on the **public subnet**
  - this subnet has a connection to the **Internet Gateway** of the VPC
- The picture completes slightly to include security via AWS IAM and VPC Peering which connects to Amazon VPC. The virtual private gateway points to on-prem data center and the Internet-Gateway points out to the public internet.

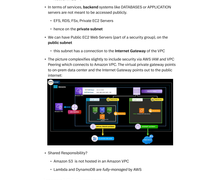

- Shared Responsibility?
  - Amazon S3 is not hosted in an Amazon VPC
  - Lambda and DynamoDB are fully-managed by AWS

Amazon VPC Components

Glance:

- CIDR Block
- Subnets
- Route Table
- DHCP Options Set
- NAT Devices
- Network ACLs
- Security Groups
- Different Types of Gateways

CIDR Block

- allows you to specify the size of your network
  - allowed block size for a VPC is between /16 to /28 netmask.
  - a netmask (subnet mask) tells you the total number of available hosts for your network.
    - /16 = 65536 IP Address
    - $/32 - 32 = 32$ $\text{hosts} = 2^{32-n} = 1$
    - /16 = 16384
    - /28 = 16
- Amazon EC2 is launched into Amazon from your CIDR block.
  - 10.0.0.0 - Network Address
  - 10.0.0.1 - VPC Router
  - 10.0.0.2 - DNS Server
  - 10.0.0.3 - Reserved for Future Use
  - 10.0.0.255 - Network Broadcast Address

Subnets

Route Table

- the **implicit router** in Amazon VPC
- controls the network traffic in your VPC through subnet routing.
- all subnets in your VPC **must** be associated with a route table
- a route table can either be the main route table OR a custom route table.
- a subnet in your VPC **can only** be associated with one route table at a time but you can associate multiple subnets with the same subnet route table.

DHCP Options Set

- a set of options that controls the automatic provisioning of IP addresses to Amazon EC2 instances and other resources
- allocates an IP address to every host, virtual machine, Amazon EC2 instance, RDS database, load balancer, or any other AWS resources in your VPC
- configures DNS, NetBios Name Server, and Network Time Protocol (NTP)

NAT Devices

- IPv4 = NAT Instance / NAT Gateway.
- Uses **Network Address Translation** (NAT)
- enable Amazon EC2 instances that are in a private subnet to connect to the public internet or other AWS services
- prevents the public internet from initiating connections with your private EC2 instances.
- works like a one-way street which means only the traffic initiated within your VPC is allowed. but not vice-versa.

NAT Instance, Gateway:

- Instance:
  - a empahsized NAT device running in an EC2 instance within your VPC
  - managed by the customer ioi
  - not highly available, not scalable.
- Gateway:
  - an advanced NAT device that is not running in your VPC
  - managed by AWS
  - highly available and scalable.
  - note that 1 NAT Gateway being routed to from 0.0.0.0/0 in an AZ public subnet can be used by EC2 instances in other AZs (private subnets in other AZ results in cheaper hourly charges, but then incurs cross AZ transfer with availability being a SPOF (single point of failure)
  - a Gateway must be in a public subnet if it wants instances in other private subnets to point 0.0.0.0/0 at it. **otherwise they will not reach the internet**
  - any subnet. whose route table points at the NAT gateway is served by it - AZ-scoped.
  - private EC2 -> private route table -> NAT Gateway -> Elastic IP -> public route table -> Internet gateway -> internet.

Network ACLs

- packet is sent from security group, and abides by it's rules first on **outbound**
- on the way back in, the Network ACL will be hit first, and then passed to the security group

Security Groups

- private subnet -> security group -> network acl -> inf gateway.

Different Types of Gateways

- Internet Gateway
- Customer Gateway
- Virtual Private Gateway
- Carrier Gateways
  - for VPCs that use AWS Wavelength to deliver ultra-low latency applications for 5G devices
  - allows incoming traffic from a carrier network in a specific location
  - allows outgoing traffic to the carrier network and to the public internet
  - only available for VPCs that contain subnets in a **Wavelength zone**
- Egress-only Internet Gateway
  - IPv6

Amazon VPC Network Architectures

Glance:

- Default VPC
- Amazon VPC
- Network ACL vs Security Group

Default VPC

- VPC Peering is **non-transitive**.
  - peering is a private wire, not a router.
- there is a default VPC in **each** AWS region
  - can immediately be used to launch Amazon EC2 instances, ELB, RDS, etc.
- good for quickly launching simple public sites / applications
- can be configured
- has an attached **Internet Gateway** by default

Amazon VPC

- can be configured to have a single public subnet
  - or public / private
    - with Hardware VPN access
    - can also drop public too.

Network ACL vs Security Group

- created by default when we launch a new VPC and on the default VPC
- acts as a virtual firewall that protects AWS resources from unauthorized traffic
- inbound & outbound rules can be set to have one IP address or a CIDR range as a source
- allows you to control the incoming and outgoing traffic to and from your subnet.

Security Group

- cannot explicitly DENY traffic
- **STATEFUL** = the security group remembers the connection.
  - so the reply is never checked against outbound rules.
  - a reply to a permitted inbound request leaves automatically. its outbound rules are never consulted for that reply.
    - they given only the connections the instance itself starts.
- ---
- tracks all the status of the incoming requests
- if a traffic is a response to a particular request, then it will be allowed automatically regardless of any rules in your Outbound Rules
- it is aware if the outgoing traffic is:
  - initiated from the EC2 instance itself
  - a response to the request that was initiated externally
- its **outbound rule** can filter:
  - an RPI call initiated by an application hosted in the EC2 instance
  - a scheduled OS Patch that is initiated by the EC2 which automatically fetches updates from a designated repository.
- no rule number
- evaluates ALL of the rules at the same time.
- applies the rules to a single EC2 instance only or to a group of AWS resources where it is associated with.
- does not use ephemeral ports!

Network ACL

- can explicitly DENY traffic
- **STATELESS**
- reflex: ask "what does the return packet look like?"
  - security group >> never ask (it remembers)
- AWS rules of explicit deny overruling explicit allow **do not apply** to Network ACL inbound rules
  - these rules just read the table from top-bottom and then a match exits out of the logic
- ---
- the inbound traffic that has already been permitted before is still subject to the rules for the outbound traffic, and vice versa.
- provides a more fine-grained control to configure both the inbound and outbound rules
- each rule has a corresponding rule number
  - evaluates rules in order starting with the lowest numbered rule
- applies the rules of all EC2 instances and other AWS resources in the subnets that it associated with.
- uses 'ephemeral ports'

V1: VPC Anatomy

- 'public' is a routing fact, not a setting.
  - there is no (public) subnet checkbox. a subnet is public iff its route table sends 0.0.0.0/0 to an internet gateway.
  - three things must hold before an instance reaches the internet, and our scenario only had the third:
    - internet gateway (IGW) attached to the VPC.
    - 0.0.0.0/0 → IGW route on that subnet's route table
    - is a public IPv4 address on the instance?

V2: Gateways and NAT

V3: Security group vs network ACL

V4: Leaving the VPC

Day 1 slides notes:

- EBS is faster than S3 and EFS. (stats retrieval)

TutorialsDojo Topic Test (VPC)

Q1. IPv4 bit. how to scale vertically and fix immediate failures?

- set up a new IPv6-only subnet with a large CIDR block and associate new subnet with the VPC then launch the instance.
  - new IPv4 not available, because that can be exhausted easily..

Q2. SAP HANA for EMF75. What to configure to allow VPC to connect through to AWS VPN?

- an internet-routable IP address (static) of the customer gateway's external interface for the on-prem network.
- by default, instances that you launch into a private virtual cloud (VPC) can't communicate with your own network.
  - you can enable access to your network from your VPC by attaching a hardware virtual private gateway to the VPC, creating a customer gateway, updating security group rules and creating an AWS-managed VPN connection.
  - although VPN is a general term, in Amazon VPC does it refers to the connection between VPC and your own network.
- a customer gateway is a physical device or software application on your side of the VPN connection.
- most importantly, we **don't need a NAT instance** to create a VPN connection.
- EIP does not attach to a Virtual Private Gateway

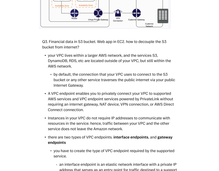

Q3. Financial data in S3 bucket. Web app in EC2. how to decouple the S3 bucket from access?

- your VPC lives within a larger AWS network, and the services S3, DynamoDB, RDS, etc are located outside of your VPC, but still within the AWS network.
  - by default, the connection that your VPC uses to connect to the S3 bucket or any other service traverses the public internet via your public Internet Gateway.
- A VPC endpoint enables you to privately connect your VPC to supported AWS services and VPC endpoint services powered by PrivateLink without requiring an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection.
- Instances in your VPC do not require IP addresses to communicate with resources in the service. Hence, traffic between your VPC and the other service does not leave the Amazon network.
- there are two types of VPC endpoints: **interface endpoints**, and **gateway endpoints**.
  - you have to create the type of VPC endpoint required by the supported service.
    - an interface endpoint is an elastic network interface with a private IP address that serves as an entry point for traffic destined to a support service.
    - a gateway endpoint is a gateway that is a target for a specified route in your route table, used for traffic destined to a supported AWS service.

Amazon Web Services Cloud
Amazon Simple Storage Service
Amazon Web Services Cloud
Amazon Simple Storage Service
Amazon Virtual Private Cloud
Amazon Virtual Private Cloud
Amazon Simple Storage Service

Day 3. Saturday.

Warm-up.

- S3 and DynamoDB hit **Gateway Endpoint**, whereas all the other AWS services hit **Interface Endpoints.**
  - hourly charge for S3 and DynamoDB are free. (Gateway)
  - Hourly Charge and Per-GB charge are none for Gateway.
    - note that Interface charges Per-hour PER Availability Zone
    - per-GB (0.01) is shared.
- Indeed it is cheaper to buy 'reserved instances' or a 'Savings plan' covering baseline EC2 instances
  - additionally, for the Heavy-use you can use 'Spot instances'
    - nightly transcoding batch.
  - commitment buys the discount on predictable load
    - interruptibility buys the discount on flexible load.
- Capacity Reservation buys capacity; Savings Plans and Reserved instances buy discount.
- every *reserved instance* -- **Standard** or **convertible**, is an EC2-only instrument.
  - it is proved to an instance type, so moving to Fargate or Lambda takes the discount to zero.
  - a **Compute Savings Plan** commits dollars per hour, not instances.
    - it applies itself automatically across EC2 (any family, size, operating system, tenancy, Region), Fargate and Lambda.
    - there's nothing pinned, so there's nothing to exchange.
  - savings plans are more flexible than Reserved Instances and reach roughly the same discount
    - yet Reserved Instances exist because we can book availability.
    - furthermore, savings plans can't be resold, transferred or cancelled.
- Standard is sold, Convertible is Converted...
  - the flexible one cannot be traded away, the tradeable one cannot be changed.
  - every other Reserved Instance property -- Availability zone, network, instance size within the same family belongs to both.
- All Upfront -> Highest savings.
  - partial upfront -> middle.
  - no upfront -> least discount, billed monthly.

Amazon EC2 Overview.

- Shared Responsibility
- can be integrated with lots of AWS Services.
- contains:
  - CPU
  - Memory (RAM)
  - Network -> VPC; Elastic IP; Elastic Network Interface (ENI); Placement Groups; Elastic Network Adapter (ENA); Elastic Fabric Adapter (EFA)
  - Disk Image (ISO) -> Amazon Machine Image (AMI)
  - SSD/HDD Storage -> Instance Store; Amazon EBS
  - Shared File Server -> EFS, FSx for Lustre, FSx for Windows File Server
  - Object Storage -> S3
  - Auto Scaling -> EC2 Auto Scaling

Instance Purchasing Options

- Rack, with rack-mounted servers, with a single server instantiating multiple EC2 instances.

On-Demand

- no interruptions
- pay by the second (minimum of 1 min) on **Linux**
- pay by the hour (minimum of 1 hour) on **Windows**

use cases:

- mission critical workloads.
- short-trem
- handling steady state
- running the **master node** and the **core nodes** of your Amazon EMR cluster

capacity reservation:

- allows reservation of EC2 capacity for a specific AZ (Availability Zone) for a period of time.
- No one-year or three-year term reservation / commitment
- good for when you need compute for a week or few months

Spot

- low supply = high price.
- lowest cost.
- spare or unused EC2 Capacity (surplus)
- automatically killed for priority customers (on-demand, etc orders)
- based on the "spot market"
  - lower prices

features:

- discounts of up to 90% compared to on-demand.
- most cost effective.
- suitable for **non-critical / infrequent jobs**

use cases:

- good for development / test environments
- apps with flexible start and end times
- interruptible workloads that can handle failures gracefully
- handling peak load / additional load onto of the usual Reserved / On-demand EC2 instances
- highly dynamic batch processing where each job is stateless

spot fleet

- collection (fleet) of Spot Instances
- can have on-demand instances

spot block

- specify a "block of time" for which the block will continuously run
- rarely interrupted compared to regular spot instances

Reserved

- 1 year / 3 years
- all upfront = full price ; partial upfront = partial price; no upfront = monthly, highest -> least discount.
- **Standard Reserved Instance**
  - can be sold in the reserved instance marketplace
  - cannot be exchanged for any other reserved instance
- **Convertible Reserved Instance**
  - cannot be sold in the Reserved Instance marketplace
  - can be exchanged for another convertible reserved instance with a different configuration, including instance family, operating system, and tenancy.
- Note:
  - both (above) can modify the attributes such as the AZ / Network
    - both can also modify the **Instance Size** using other sizes within the same instance family.
    - both require a fixed 1-year / 3-year commitment.

Scope:

- Standard = Regional (I think)
  - encompasses Availability Zones
- Convertible = Zonal
  - encompasses just an availability zone.

Use Cases: (for either Standard or Convertible)

- running non-interruptible workloads for a one-year or three-year time frame
- workloads with predictable capacity and uptime requirements
- hosting the application servers of your production environment
- for processing the steady-state load or the baseline capacity of workloads
- batch jobs that **cannot** be interrupted once started
- for consuming SQS queue messages in which the application should continually process messages without any downtime
- running the Master node or core nodes of your EMR (Elastic MapReduce Cluster) -- cheaper than On-Demand Instances

Dedicated

Hosts

- a single, physical rack-mounted server. a.k.a a host.
- per-core = CPU Core
- for cases when the existing server-bound software licenses must be used by customers
- to comply with per-core software license requirements
- for compliance and software licensing requirements mandating that a workload must be hosted on a physical server
- cost analysis that supports physical isolation of a customer workload
- Launching Windows Server, SQL Server, SUSE Linux, ... and other software licenses bound to particular VMs, sockets or physical CPU cores

Instances

- Regular virtual machines that run in a virtual private cloud (VPC) on hardware that's dedicated to a single customer
- Dedicated Instances that belong to different AWS accounts are physically isolated at a hardware level
- may share hardware with other EC2 instances if the instances are:
  - in the same AWS account
  - not a type of dedicated instance
- allows you to launch Dedicated Spot Instances, Dedicated On-Demand Instances, or Dedicated Reserved Instances

Savings Plan

Features:

- save on EC2, Fargate, Lambda
- provides discounts in exchange for a commitment to a consistent usage amount that is measured in dollars per hour for a one or three-year term
- can be purchased from aws account, management account, member account

vs reserved instances:

- both require a fixed one-year or three-year commitment
- both provides Billing Discounts
- note that:
  - savings plans
    - are based on consistent amount of compute usage
    - provides flexibility to use a more suitable compute option at low prices without any exchanges or modification
  - whilst, reserved is
    - based on a specific instance type or instance size
    - must exchange or modify the reserved instance to suit your current requirements

Capacity Reservation

features:

- allows you to reserve capacity for EC2 instances in a specific AZ
- independent of billing discounts offered by **Savings plan** or **regional Reserved Instances**
- works like a zonal reserved instance
- no 1-year / 3-year commitment
- you can reserve a particular AZ only (zonal), no regional reservation in scope
- can be applied to on-demand EC2 instances

REQUIREMENTS

- AZ
- number of EC2 instances
- instance attributes.

EC2 Instance Types

- cpu, ram, storage, graphics, network ALL CHOICES
  - cpu options: intel, amd, graviton.
- instance family
  - instance sizes:
    - nano, micro, small, medium, large, xlarge, metal

Instance Categories

- General Purpose: Mac, T*, M*, A*
- Compute Optimised: C*
- Memory Optimised: R*, X*, Z*, U*
- Storage Optimised: I*, D*, H*
- Accelerated Optimised: P*, Inf*, G*, F*
- Others

Naming Conventions.

Instance Naming Conventions:

- m6 => m = 'Instance Family', 6 = Generation.
  - m4 & below = previous gen
  - m5 = 5th gen
  - m6 = 6th gen
  - *m7 = = next gen*
- **a = AMD
- **g = AWS Graviton
- t3a => 'a' = AMD
- m6g => 'g' AWS Graviton
- t3, m5, r5 => intel
- ***d => has a local NVMe-based SSD storage
- ***n => has enhanced networking capabilities
- * n => burstable performance instances
  - provides a baseline level of CPU performance with the ability to burst above the baseline
  - ability to burst is governed by CPU credits
    - accrued when the instance is idle.
    - a sort of 'vertical scaling' since it temporarily provides higher CPU performance over the maximum CPU capacity of the instance
- metal.
  - grants direct access to the CPU and memory resources of the underlying server
  - doesn't have pre-installed KVM, Xen or AWS Nitro Hypervisor that other EC2 instances use
  - allows you to fully access the CPU, Storage, and Networking bandwidth of the underlying server.
  - allows customers to run their own hypervisor or virtualisation secured containers such as Clear Linux Containers.

Quiz EC2

Q1. CloudWatch

Q2. Create a role in IAM and assign it to the EC2 instance

Q3. Reserved won't work because it's contract 1, 3 years.

Q4. Egress-only Internet Gateway to the VPC to allow outbound ipv6-communication to the internet. Use AWS Network Firewall to set up the required rules for traffic inspection nad traffic filtering.

Q5. **Memory Utilisation** is non-default in CloudWatch

Q6. on-demand instances cannot reserve compute capacity at all. on-demand capacity reservations can be good. Dedicated Hosts will waste time during off-compute hours.

Q7. Multiple EC2 instances in multiple AZ's are better for fault tolerance.

Q8. There is a VCPU-based On-Demand Instance limit per region which is why subsequent requests failed. Just submit the limit increase form to AWS and retry the failed requests once approved.

Day 4. Sunday.

Probably more lab work today.

Review:

- CloudWatch publishes for free, exactly what the hypervisor can see from outside the instance:
  - CPU utilisation, network in/out and packets, and block-device read/write operations and bytes
  - it cannot see inside the guest operating system, so everything that only the OS knows needs the agent: memory Used, disk space used, swap used, page-file utilisation, log files.
  - "Disk read activity" is the hypervisor counting blocks crossing the boundary. "Disk spaced used" is a filesystem concept that exists only inside the guest.
    - if a metric requires knowing what the OS thinks, it needs the agent.
- the On-Demand quota is counted in **virtual CPUs per Region**
  - running instances + request can exceed it.
    - it can be raised by request

Amazon EC2

AMI - Amazon Machine Image

- regional in scope
- *you can* ***copy*** *your AMI to another AWS region*
- *you can also copy your AMI to another AWS account*
  - *i don't think you can copy your EC2 instance to another account*
- **Volume Snapshots**
  - EBS Snapshots
  - Template for the root volume.
- **Block Device Mapping**
  - amazon EBS Volumes mapping
- **Launch Permissions**
  - public, explicit, implicit

AWS Marketplace

- ParaVirtual = PV.
  - uses special boot loader called PV-GRUP
  - no support for special hardware extensions such as enhanced networking or GPU processing
- HVM = Hardware Virtual Machine
  - executes the master boot record of the root block device of your image.
  - indeed support for several hardware extensions such as enhanced networking or GPU processing

Instance User Data

- must be in a base64-encoded format
- limited to 16KB only when in raw form
- accessible from the Instance Metadata using this URI:
  - https://169.254.169.254/latest/user-data
- only run once upon the first EC2 Instance Launch
- Modifying the User Data and restarting the instance won't affect the initial User Data

Instance Metadata

- curl the endpoint.

EC2 Networking

- public internet, private network in aws
- physical networking devices, virtual network interface card
- ip addressing, elastic network interface
  - elastic ip address, enhanced networking, EFA (Elastic Fabric Adapter)
  - Elastic Network Interface
    - primary private Ipv4 address
    - secondary too
    - on elastic ip / private ipv4
    - one public ipv4
    - one OR more ipv6
    - one or more **security groups**
    - MAC addresses (media access control)
    - source-destination check flag
    - custom description
- CIDR = Classless Inter-Domain Routing
  - a method for allocating IP addresses
  - also used for IP routing
- **Enhanced Networking**
  - based on the network adapter drivers of the underlying physical host
  - the network adapter drivers can be:
    - intel network adapter virtual function driver
    - aws-built custom-based network adapter driver called ENA (Elastic Network Adapter)
    - Network drivers provided by AWS or other companies
  - similar to the 'driver' that comes with Windows / Linux software.
  - uses single root I/O virtualisation or SR-IOV
  - provides higher I/O performance and lower CPU utilisation than the traditional virtualisation techniques
  - controlled by network drivers (software)
  - provides:
    - higher bandwidth
    - consistent lower inter-instance latencies
    - higher packet per second performance (PPS)
  - EFA (Elastic Fabric Adapter)
    - like ENI (Elastic Network Interface) but with additional capabilities
    - can directly communicate to the network interface hardware without passing through the Linux Kernel — aka OS-bypass
    - provides low-latency and reliable transport functionality to your virtual machines
    - accelerates the networking capabilities of your HPC workloads
    - enhances inter-instance communication

Security Groups and Network ACLs

- Default Network ACL allows all inbound and outbound traffic by default.
- Custom Network ACL denies all inbound and outbound traffic by default
- 1 EC2 instance can have one or more security groups
  - Cannot have an explicity DENY rule (unlike Network ACL)
  - aside from EC2, can also be attached to RDS, ElasticCache, etc.
- can't apply either of these to Amazon S3 buckets.
  - they don't provide enough protection against Cross-Site scripting or SQL Injection attacks
  - also ineffective at geographic match conditions or blocking certain countries.

Placement Groups

- Cluster
  - provide low-latency network performance and high network throughput
  - logical group. host / rack networking.
- Partition
  - commonly used on large distributed and replicated workloads
    - Hadoop, Cassandra, Kafka
  - **separates groups**
- Spread
  - reduces correlated failures and improves availability.
  - **separates instances**

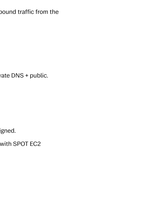

EC2 AutoScaling

Elasticity

Scaling Types:

- Vertical Scaling
- Horizontal Scaling

Auto-Scaling Types:

- Simple Scaling
- Step Scaling
- Target Tracking
- Scheduled Scaling

Lifecycle Hooks

- It is **not** the case that an instance in the wait state is 'InService' and receiving traffic
  - because, if that were true, then the hook would solving nothing and errors in AZ's stem would still happen.
  - 'Pending:Wait', is precisely not 'InService'

Slides

- Ch6. Amazon EC2.
  - AMI - Amazon Machine Image
  - Instance User Data
  - Instance Metadata
  - EC2 Networking
  - Security Groups and Network ACLs
  - Placement Groups
- Ch7. EC2 Auto Scaling
  - Auto Scaling Overview
  - Scaling Types
  - Lifecycle Hooks

Labs

- A public subnet is a subnet that has a route to an internet gateway in its route table.
  - this links the VPC to the internet as well as other publicly accessible AWS services such as Amazon S3 and Amazon DynamoDB
- A private subnet is one that has a route table with no connection to an internet gateway.
  - *Resources inside a private subnet do not allow inbound traffic from the internet.*

Launching an EC2

- straight forward. just press launch.
- choose an AMI and deploy.
- manage your security group rules and you'll get a private DNS + public.
- create a KeyPair with .pem file.

Launching an EC2 Spot Instance

- EC2 -> Spot Requests
- 10 means high probability, 1 means you won't get assigned.
- recall that they can pull the rug from you at any time with SPOT EC2 instances.

Vertically Scaling an EC2 instance

- any kind of bump in specs.
- you'll have to stop the instance.
  - 'instance state' drop-down -> 'Stop Instance'
  - 'Actions' dropdown -> 'Instance Settings' -> 'Change Instance Type'
    - you can control EBS optimisation. whatever that is
- start the instance. done!

I'm getting a little confused.

Enhanced Networking
↳ driver
↳ single root IOV
SRIOV I/O virtualisation

Elastic Fibre Adapter
↳ device
OS-Bypass

The sensitive is if the OS matters here, the OS stands in digitally complex HPC, it's EFA.
If it goes to throughput, bandwidth, or packets
↳ think Enhanced Networking

Not alternatives! EFA is ENI with extra capabilities.

Day 5, page 5. Monday.

Yes, a splash of frizz today; casual. Furthermore, I did the first set.

Schedule.

- Set 1. 53%
- Set 1 Review.
- 99 Slides
- Anki.

Review

Q1. Which service/s allow SMB capability?

Q3. Amazon Data Firehose. What does Redshift do again?

Q4. FSx for Windows/Linux vs. EFS?

Q5. AWS Directory Service AD Connector? Also, what is Service Simple AD?

- AWS Directory Service allows Amazon Cloud Directory and Microsoft Active Directory (AD) to be used with other AWS Services.
- LDAP = Lightweight Directory Access Protocol.
- IAM Groups are just a collection of IAM Users. Groups let you specify permissions for multiple users, which can make it easier to manage the permissions for those users.

Q6. When to use S3 server-side encryption vs. when to use client-side?

- Client-side encryption is the act of encrypting data **before sending it to Amazon S3**. To enable client-side encryption we could:
  - use an AWS KMS key
  - use a client-side master key.

Q7. Note that IOPS SSD is for hot storage. So you got this wrong :(

Q8. Launching an AWS Transfer for SFTP endpoint vs configuring?

- ELastic File System (EFS) vs S3 Bucket?

- 30-day retention policy must be primarily configured on the Amazon S3 bucket. there's no retention policy option on AWS Transfer for SFTP
- EFS lifecycle management doesn't delete objects. it can only transition files in and out of the 'Infrequent Access' tier.

Q9. Amazon Kendra vs. Polly Vs. Macie, vs Fraud Detector

- Macie is an ML-powered security service that helps you prevent data loss by automatically discovering, classifying and protecting sensitive data stored in S3.
  - PII, IP.
  - generates **policy findings** and **sensitive data findings**
- Amazon Polly is simply a service that turns text into lifelike speech.
- Amazon Kendra is just an enterprise search service that allows developers to add search capabilities to their applications.
- Amazon Fraud Detector is a fully managed service for identifying potential fraudulent activities and catching onlie fraud faster -- it does not check any S3 data containing PII

Q12. IAM, STS?

Q15. RDS Read Replica, or DynamoDB Read replica?

Q16. Canary-Release; blue-green deployment strategy

Q17. Cross-Region Replication (CRR)

Q19. AWS Artifact, Inspector

Q20. Recovery Point Objective (RPO) - 1 second.
Recovery Time Objective (RTO) - 1 minute. What are these?

Q21. What's the relationship between Aurora and RDS?

Q23. Signed URL, signed cookies, Match Viewer

Q25. Datasync or site-to-site VPN?
How much throughput can DataSync handle?

Q27. Amazon EventBridge

Q28. What exactly is an ECS Cluster?

Q30. What is S3-Intelligent Tiering?

Q34. AWS Glue or AWS Lambda?

Q38. target tracking scaling vs suspend & resume vs. simple vs. scheduled

- which are dynamic?

Q39. can you use AWS EFS with Microsoft Server? (EC2)

Q40. Partition keys? What *exactly* is DynamoDB?

Q41. Retention lock, legal hold? Compliance vs governance.

Q42. Is DynamoDB better with schema changes or Aurora?

Q44. Tape gateway, file gateway

Q49. How does the Auto-scaling group punt EC2 instances?

- least users, first launched, etc?

Q50. SAML? Web Identity Federation?

Q54. What is Amazon Inspector?

Q55. Redis Replication Groups?

Q56. Enhanced Monitoring in RDS? Is there such thing?

Q58. AWS Lake Formation. What is it, and what does it do?

Q60. AWS KMS. Key management Service.

- what can it do, and what can't it do?

Q65. AWS VPN CloudHub?

Day 6.I woke up before IIAM.

Slides

Glance:

- EBS
  - Overview
  - Volume Types
- ELB
  - Overview
  - Types: ALB, NLB, GWLB
- S3
  - Overview
  - Storage Classes
  - Event Notifications.

EBS

Overview

- Elastic Block Store.
- type of block storage like EC2 Instance Store (which I didn't know was a block storage either)
- data is more persistent and will not get lost even if the EC2 instance is stopped/rebooted or terminated.
- Zonal in scope — i.e. exists only in a single Availability Zone
  - can be attached to any EC2 instance in the **same AZ only**.
- can be encrypted at rest with AWS KMS
- you can attach one or more EBS volumes in a single EC2 instance.
- suitable for lots of workloads
- high throughput, low latency
- the EC2 + EBS volumes are logically together and are both located with a single AZ — this significantly reduces latency.
- mainly operates at the hardware level.

Raid

- Redundant Array of Independent Disks
- RAID 0 -> Splits data evenly across two or more disks, without parity of information. SPEED IS THE intended goal
  - **striped** multiple volumes together
  - provides greater I/O performance
  - divides a body of data into blocks and then spreads the data blocks across multiple storage devices
  - suitable of I/O performance is priority
- RAID 1 -> An exact copy of a set of data on two or more disks. **Performance / Reliability** is more important.
  - **mirrors** two or more volumes together
  - provides on-instance redundancy
  - duplicate data to provide more durability and availability
  - suitable if data redundancy is your focus
- SSD, HDD:
  - SSD fast, IOPS = Input/Out operations per second
  - Dominant Performance Attribute: HDD slow -> for data archiving, backups. Throughput
    - SSD can be used as Boot Volume for EC2, HDD cannot

Snapshots:

- incremental backup that uses S3 internally.
- only saves data blocks that have changed since last snapshot
- we can restore, obviously.
- allows copying of EBS volume to another AWS region
- can be used to encrypt an unencrypted EBS volume.
- automate the creation, retention, and deletion of EBS snapshots and EBS-backed AMIs using the Amazon DLM (data lifecycle manager) service

Volume Types

- Root EBS Volume.
- Other Data Volumes
- **SSD**:
  - suitable for transactional workloads
  - for various types of applications and systems with frequent read/write operations with small I/O sizes
  - performance attribute: IOPS
  - General Purpose SSD or Provisioned IOPS SSD
    - Provisioned is good for sub-ms workloads on small filesystems.
- **HDD**:
  - large streaming workloads
  - large, sequential I/O operations
  - performance attribute: throughput
  - Throughput Optimized HDD or Cold HDD

Anti-patterns

- temp storage should be EC2 Instance Store
- if you need POSIX compliant hierarchical directory structure use Amazon EFS
- concurrent file modifications should use Amazon EFS / Amazon FSx
- store static data in S3 — it's more cost effective.

ELB (Elastic Load Balancing)

Overview

- "the distribution of traffic to underlying resources"
-

Types: ALB, NLB, GWLB

- ALB = Application Load Balancer.
- NLB = Network Load Balancer
- GWLB = Gateway Load Balancer
- CLB = Classic Load Balancer

types	ALB	NLB	GWLB	CLB
routing algorithm	round robin, least outstanding requests (LOR)	Flow Hash	IP Listener Routing that leverages on GENEVE protocol	Round Robin Least Outstanding Requests
protocol listeners	HTTP/HTTPS gRPC	TCP/UDP TLS	IP	HTTP/HTTPS TCP SSL/ TLS
use cases	web apps, microservices & containers	handling millions of requests per second while maintaining ultra-low latencies	running third-party virtual appliances in AWS	for legacy apps in AWS ; for implementing custom security policies (CSP?) and TCP passthrough configuration.

S3

Overview

- object storage
- Simple Storage Service
- highly durable, available and scalable storage service.

Storage Classes

- S3 Standard
  - for frequently accessed data
- S3 Intelligent-Tiering
  - for changing or unknown access patterns
- S3 Standard-IA (Infrequent Access)
- S3 One Zone-IA (Infrequent Access) for storing long-lived, yet less frequently accessed data
- S3 Glacier
- S3 Glacier Deep Archive for low-cost, long-term storage and data archiving

LifeCycle Policy:

S3 Bucket
Object
Object
Object

Day 7. Wasted away.

Wam looks promising though. Dishes got washed, laundry is done.

- Multi-AZ RDS keeps the endpoint stable by rewriting a CNAME; Aurora keeps it stable with a cluster endpoint that resolves to the current primary.
- An LSI lives inside each partition — it shares the partition key and only varies the sort key. Index and data sit in the same place, so a read can be answered there and strong consistency is available; and because it is part of that partition's physical layout, it must exist when the table is built.
- A GSI is a separate, table-wide structure with its own partition key, so a query fans across everything and the index is maintained asynchronously — which is exactly why it can promise only eventual consistency, and exactly why it can be bolted on to a live table.
- Note that a Standard Reserved Instance **locks** the instance family and cannot follow the workload to Fargate or Lamda.
- UDP and static IPs each rule out CloudFront and Route 53 on their own.
  - only HTTP traffic (non UDP, TCP, gaming, IoT, VoIP) goes to CloudFront
  - CloudFront gives you a domain name, Route 53 gives you DNS, Global ACcelerator gives you anycast IPs
- Amazon Data Firehose. Not AWS Glue...
  - no code for Data Firehose.
  - Streams you consume. Firehose delivers. Glue transforms. Athena asks.

	Local Secondary Index	**Global Secondary Index**
Scope of a query	one partition (localised)	all partitions of the table
Consistency	eventual **or strong**, your choice	**eventual only**
When creatable	**only with the base table**	any time; deletable any time

	D1 · secure	**D2 · resilient**	**D3 · performing**	**D4 · cost**
Storage	encryption, Object Lock	backups, replication	S3 vs EFS vs EBS	tiers, lifecycle
Compute	roles on instances	Auto Scaling, health	Batch, EMR, Fargate, sizing	Spot vs Reserved vs Savings Plans
Database	KMS, at-rest encryption	Multi-AZ, read replicas	ElastiCache, IOPS, Aurora vs DynamoDB	DynamoDB vs RDS, engine choice
Network	security groups, NACLs, WAF	Route 53 failover	CloudFront vs Global Accelerator	NAT cost, VPC endpoints

Sick. Signing off. 10:22PM.

Day 8.

	Endpoint	Points at
Cluster	writer	whichever instance is **currently primary**
Reader	load-balances across **all Aurora Replicas**	ready, membership auto-updates
Custom	**a subset you define**	e.g. send analytics to your two largest replicas
Instance	**one named instance**	diagnosis, breaks on failover

- okay, but truthfully you can't just keep sale orphing the table. it's not **open book**
- LSI is local to a partition. Same partition key, different sort key, stored beside the item.
  - GSI indexed can be added not he table after the table already exists and holds data.
  - its partition key may also be an attribute other than the table's partition key.
  - On the other hand, LSI draws on the base table's provisioned throughput rather than having it's own
    - also a query against it is confined to a single value of the table's partition key.
    -

	LSI	GSI
Partition key	**must** be the table's one partition key value	any attribute whole table
Sort key scope	Consistency	strong available **eventual only**
Throughput	shares the table's	**its own**
When creatable	**table creation only**	any time

Slides.

1. Intro
2. AWS Overview
3. Services Overview
4. IAM
5. VPC
6. EC2
7. EC2 Auto Scaling
8. EBS
9. ELB
10. S3
11. **RDS and Aurora**
12. **DynamoDB**
13. **Application Integration**
14. AWS Analytic
15. Serverless Computing
16. Amazon Route 53
17. CloudFront
18. Comparisons

II. RDS & Aurora

RDS Overview, pg 773-836

- relational database service
- managed by both client and AWS
- allows various database engines to be run
  - SQL Server (MSFT)
  - mySql
  - MariaDB, PostgreSQL, Oracle, Amazon Aurora
- can be deployed using CloudFormation, CLI, RDS API or Management Console.
- allows configuration of underlying EC2 instance
- Comparisons against *Self-Hosted database* on EC2 instance, vs. RDS db:
  - Amazon RDS *cannot* be accessed via SSH, RDP
  - modify database configuration in RDS via
    - parameter group
    - options group
- OLTP = Online Transaction Processing Applications
- ACID = Atomic, Consistent, Isolated, Durable
- RDS proxy automatically connects application to a new DB instance whil preserving its application connections
  - minimises downtime by instantly routing the incoming requests directly to the new database instance

Event Notification

- Source Type: Instances, Security Groups, Parameter Groups, Snapshots, Clusters, Cluster Snapshots
  - Event Categories
  - Target Type
- Fanout Event Notifications

- VPC spans all the AZs in a region.
- subnet is locked to one AZ.
- VPC cannot span regions. (it's a regional construct)
- crossing regions means VPC peering.
- Subnet
- Gateway
- VAE Route 53, CloudFront, S3 bucket namespace
- Regional
- VPC, EBS snapshots, security groups, S3 buckets, DynamoDB tables
- AZ
- subnets, EBS volumes, EC2 instances.

Multi-AZ Deployments

- Synchronous Replication. Standby Replica.
  - Two-Way
  - can't be deployed to another AWS region
- Asynchronous Replication: Read Replica
  - one way
- Multi-AZ db will have minimal disruption during hardware maintainance, unlike Single-AZ db.
  - multi has about 1 minute unavailability.
  - CNAME is automatically altered to point to the newly promoted Standby Instance
- can only provide high availability in **a single AWS region only**
  - cannot deploy a **standby replica** to another aws region.
- cannot provide RPO of 1 second and RTO of 1 minute.
  - if we have this requirement, we'd need **Aurora Global Databases.**

Read Replicas

- can provide cross-region db replication for multi-region disaster recovery (which a Multi-AZ deployment configuration can't provide)
- Standby Replica cannot be seen in the RDS Console as a separate DB instance.
  - DB endpoint is same as primary DB instance
  - Read Replica can accept live traffic
    - can be seen in the RDS Console as a separate DB instance.
    - the DB Endpoint is different from the primary DB instance.
- just a regular db with read-only configuration.
- can be launched on the same AWS region of primary DB
  - or on a different AWS region.
- does not allow modification of its config files
  - parameter group reqd.
- useful for:
  - db sharding, implementing failure recovery
  - DDL operations (Data Definition Language)
  - primary AWS Region outage.
- limitations:
  - cannot directly create an encrypted RR from an unencrypted instance.
    - can be made from encrypted primary instance though.
- allows use of a custom encryption key
- good for offloading read requests / reporting
  - offloads primary for write ops.
- used to primarily improve scalability
- no automatic failover

Aurora Overview, pg 838-850

- fully managed db service and also a type of database engine within Amazon RDS
- scales automatically, performs faster and costs lower
- relational database compatible with mySQL, postgreSQL
- usually deployed as a database cluster
  - which contains **one primary writer / reader**
    - and multiple replicas. readers
- Cluster types:
  - Single-master
  - multi-master.
- Standalone Type
  - single primary DB instance with no replica
- can scale computing components and storage automatically
- db cluster typically lags behind the primary instance by a few ms
- provides less than 1 second of read replication latency for Aurora Replicas in the same or different AWS region.

Endpoints:

- Cluster endpoint
- Reader endpoint
- Custom endpoint
- Instance endpoint

Serverless:

- sporadic usage workloads or for unpredictable usage.
- pay for it per second.
- more cost effective than RDS or Aurora DB
- good for migrating legacy apps
- good for moving from monolithic architecture to microservices

Aurora Global Databases

- for globally distributed applications
- allows a single Aurora DB to span multiple AWS regions
- offers faster physical replication between Aurora clusters
- eliminates the need to manually create cross-region Aurora Replicas yourself
- RPO = 1 second
  - RTO < 1 minute

DynamoDB, pg 851-882

Overview

- recall that a *Relational Database* is for applications with well-defined schemas that do not change too often
  - have hundreds of thousands of tables, multiple table joins, foreign keys, complex SQL queries, relationships with other tables
  - ACID properties
  - perfect for transactional workloads
- on the other hand, we have **NoSQL Database** which has a flexible schema and changes often.
  - one table.
  - can scale better.
- DynamoDB is a fully managed NoSQL Database
  - highly scalable.
  - single-digit ms performance
  - serverless
  - least amount of operational overhead compared to other db types
  - capable of auto scaling read and write capacity
  - can be queried by simple key-value requests via API
  - instead of a Primary Key (relational), we have Primary / Partition key.
  - it seems like the 'secondary index' and 'global secondary index' are concepts for the NoSQL db.
- Local Secondary Index
  - queries data over a single partition only
  - supports both eventual consistency or strong consistency
  - can only be added at the same time that you create the base table.
- Global Secondary Index
  - queries data across all partitions of the entire table
  - only supports eventual consistency
  - can be added or deleted at any time

Features

- Streams:
  - can associate with Lambda
  - can also integrate with Kinesea Data Streams
- TTL:
  - automatically expire the items
  - Time To Live
  - reduces the number of obsolete data
    - lower costs
- Transactions:
  - ACID properties for transactional workloads
  - all-or-nothing change
  - atomic operation

DAX (DynamoDB Accelerator)

- in-memory cache
- DAX cluster
- microseconds not just milliseconds
  - significantly reduces the response times

Scaling

- measured in terms of Read Capacity Unit (RCU)
  - and Write Capacity Unit WCU
- *Provisioned Capacity Mode*
  - predictable traffic
  - manually set RCU, WCU
  - *auto-scaling*
- *On-Demand Capacity Mode*
  - inconsistent traffic
  - sharp spikes
  - no manual auto-scaling setting
  - variable traffic

Security

- protects both in transit and at rest.
- fully encrypted by default
- VPC Gateway Endpoint to avoid API calls from the private EC2 to DynamoDB to not traverse the public Internet.
  - need to add a new route table entry too.

Backups

- Point-in-time Recovery (PITR)
  - automated backup process
  - enables continuous backups to your table
  - additional costs
- On-Demand Backup and Restor
  - Manual backup process
  - no continuous backups
  - can only restore to a particular backup you've taken
  - cost-effective yet limited backup option feature for data

Core Components

- Table
- Primary Key, aka partition key.
  - Simple
    - partition key
  - Composite
    - partition key + sort key.
- Item (the row of the table)
- Attribute (col of table)
- Secondary Index
  - make queries run faster
  - provides more flexibility and performance improvement
  - supports advanced queries to access stored data faster
  - allows query the data using an alternate key other than the primary key

Application Integration pg. 884-912

- empowers migration from **Monolithic Architecture** to **Distributed Architecture**
  - **Monolithic**
    - tightly coupled
  - **Distributed Architecture**
    - loosely coupled
  - SQS, MQ, SNS, EventBridge, AppSync, API Gateway, Step Functions

SQS

- decouple tightly-coupled architecture
- process workloads async
- FIFO message queue
  - can be HTTP or API
    - fetching messages for processing is called polling
- fully-managed message queue
- for workloads with long-running requests
- assists in scaling compute resources
- **TYPES** also standard (opposition to FIFO)
  - delivery is 'at least once'
    - FIFO is 'exactly once'
  - ordering is 'best effort'
    - FIFO is 'preserves exact order'
  - throughput 'high'
    - FIFO 'low'
- **SETTINGS:**
  - visibility timeout
  - message retention period
  - delivery delay
  - wait time
  - access policy
- **SECURITY:**
  - Encryption
    - Data in-Transit
    - Data At-Rest
- SQS < (AWS Lambda ; SNS; EC2 Auto Scaling; S3 ; ECS & EKS)

Standard (?)

EFO. Dead letter queue
Delay queue
Temporary queue

SNS

- publishers, subscribers.
- Standard, Fifo
- Encryption is similar: Data In-Transit; Data At-Rest
- Features:
  - Message Filtering; Fanout, Durability; Encryption, Archiving
- DLQ = Dead Letter Queue

- 1.4.304 * Amazon Credential Manager
  - EC2 = Elastic Compute Container?
  - EC2 = Elastic Container Service?
- S3: Encryption on Direct Connect
- S3: Prevent/block/remove Video Streams, Metaphotos: App/Container. Amazon Elastic Transcoder?
- EFS: Encryption?
- S3: Does RDS have 1 Read Replica in Primary?
- S3? Redshift? Integrated Write to, Serverless Compute?
  - assign a role to the EC2 instance OR assign a user to each EC2 instance?
- S3: CloudFront vs Global Accelerator
- EFS: Elastic IP Address? what even is this. Have we covered it?
- EFS: AWS System Manager OpsCenter??
- EFS: AWS SAMT
- EFS: Best protocol for accessing other AWS services from EC2 instances?
- EFS: Routing policy: ~~~~Differences between Geolocation, Geoproximity, Weighted and Latency Routing
- EFS? What does the NAT gateway do?
- EFS: Glacier Flexible Retrieval vs. Glacier Deep Archive
  - S3 vs Storage Gateway
- EFS: EventCache vs CloudFront?
- EFS: RDS Trusted Advisor? CloudFront?
- EFS: Is Simple Email Service real?
  - CloudwatchS?
  - BIND server?
- S3/4: Is Lambda function cheaper than on-demand EC2 instances?
- EFS: AWS Backup vs Amazon DLM?
- EFS: All options seem feasible. EBS vs S3?
- EFS? Document model means SQL or noSQL?
- EFS: Fargate vs EC2 launch type. Open source?
- EFS: Cost Explorer vs AWS Budget
- EFS: What's an elastic IP address?
  - Network Load Balancer vs Application Load Balancer
- EFS: What is an Outposts server?
- EFS: Can you modify EBS configurations live?
- EFS: WAF is best for Cross-Scripting attacks right?
- EFS: Can you encrypt the data on an S3 bucket?
  - AES-128 vs AES-256?
- EFS: Does Route 53 have a failover option?
- EFS: CloudTrail vs. Inspector?
- EFS? Auto-scaling cost-doesn't what is that?
- EFS:AWS Data Transfer Terminal
- EFS: Is CloudWatch installed on EC2 by default?
  - detailed monitoring: does it show swap?
  - when do I need cloudwatch?
- EFS: Scheme Conversion Tool? what is it, how does it work?
- EFS? ELB vs Route 53 with weighted routing
- EFS: Is HSM mandatory exclusive with AWS organizations?
- EFS: What is SAMT?

AWS gives bandwidth when because I am a person:

AWS = AWS Cert Certificate Manager

EC2 = Elastic Compute Cloud. (what ordinal)

Redshift = data warehouse
↳ columnar, petabyte-scale.
↳ distribution of an ETL job.

EC2 is not serverless

EC2 is open source

Fargate: serverless compute engine, underneath EFS / EC2 architecture

Fargate

EFS / EC2

outposts: subset that works installed in your own data center (s/o) workloads that must stay on-prem (5 delivery / data residency)

Lambda is cheaper than EC2 for bursty, short, event-driven work.
↳ (same for 24/7 workloads)

EFS = Elastic FS ← a shared 2D/3 workloads, attached to your server, file-accessible between instances.
Not appropriate on NLB (Network Load Balancer)

S3 Storage Lens: improve replicas!
↳ Cannot be attached to ALB only ALB

SQS = Serverless Application Model.
↳ a deployment framework → Cloudformation extension (and a runtime service) for serverless apps

SNS: I queue consumes pull. one consumer takes and manages

SNS: push/pull. published to many subscribers at once.
↳ targets include SQS, Lambda, HTTP/S, email & SMS.

SQS: Simple Email Service. Real!
↳ EC2 and Transactional email to customers and accounts - not spam alerts customers
↳ Verify operations team's job SNS only

SNS = Simple Notification Service. Stop flooding in the monitors replacement

CloudFrontEdge: 1 Brand by with content-based routing on the event body. Small product services, and even subdomains.

NFS updating

EBS/EFS open-source tools server run on EC2

Storage Transfer Migration

Storage Gateway. Hybrid. 3 types. Tape, File, Volume

Data transfer terminal. Boxes to ship and use their bandwidth. Integrated Work Partners

Cluster Placement Optimal. Maximizes (to lower), change.

Cluster Spread Placement. Maximizes to Minimum. Compact

DLM (Data Lifecycle Manager). Policy Automation. EBS snapshot create instead (database or no underlying cost.

NAT Gateway: Centralized outbound traffic proxy provides AZ-resilient and cross-region routing

↳ Check 3 "Easy Backend" or "Serverless Management" then go from EC2 and Backup

Networking & Content Delivery

- S3 + Conditions = the cheapest way to host static site
- Direct Accelerator Connect front on S3 hosted to centralized AWS backbone instead. Only UDP/TCP
- Configurable in console on Cloudwatch → NAT → CIDR?
- VPC can tether on Elastic IP
- Lambda@Edge: for Cloudfront origin handling
  - ↳ Quick "modify the request"

Networking Performance and Security

- Bandwidth tests for quick improvement
- Cloudfront will tell you when slow it
- Instance/Cluster gives best problem recommendations on VPC security limits. Cannot duplicate a role.
- Systems Manager Opscenter → a central where operations and DevOps and manage operational work items
  - ↳ Ask prioritizing, and governance.
- Control Tower → spell "easy accounts with pre-approved baseline, least access"
- Connectivity → Transit Gateway: connects of logs for auditing/ Regulated audit
- Inspector → Vulnerability assessment Connectivity is explicitly
- Reachability → Configuration that not cause of a Outage
- Macie → Discovery and identifies problems in S3
  - ↳ Classifies & the personal data!
- WAF, Web Application Firewall
  - ↳ good for cross-scripting, SQL and rate-based Shield too.
- Firewall Manager: Administrates WAF and Shield across many accounts.

Machine Learning

Rekognition → images and video. Logo, faces, text, moderation, celebrities

Kendra → intelligent enterprise search.

Fraud Detector → Detect payment fraud, fake accounts.

Comprehend → Search docs. Finds Clusters, Sentiment, Sensitive data, Detect language.

SageMaker → Build & train your own model

Lex → Natural-language: conversations existing AWS/Alexa apps ability to integrate with AWS...

Amazon Service for Monitoring: Cloudwatch and Cloudtrail monitoring actions and changes, video monitoring actions and changes.

Test Review. Questions Answered.

Q1

A: AWS-KMS serves as records of all

Q2: Assign an IAM user to an EC2 instance is never

A: Assign role to the instance. correct.
↳ always refer to the instance
↳ anything long credentials underneath and rotation them.

Q3: never store credentials on an instance.

Q4: S3 → Cloudfront

Q5: Encrypt in S3 SSE.

Q6: AES-256 doesn't exist in S3 and S3 always wrong
↳ You can encrypt an S3 bucket

Q7: Route 53 always runs a failover option (redundant personally)

Q8: Security Analysis: a Certificate manager - ~~~~ the good Publicly and automatically renew before the expiring
Security Analysis for Cloud Advisor.
↳ 1 2 set a root password for the instance.

Q9: Cloudwatch, not installed by default.
Detailed monitoring you only even enhance, it only changes the granularity of them!

Q10: Signature group mismatch. IAM controls identities and permissions inside an account. IAM identity ~~~~ controls who does EC2, and who gives one instance access to of them.
↳ Not actually available.

Random Questions

- Compromised policy (Above S3) → when one wants to control bucket level. Enable them to publish from a region.
- Cloudwatch → for Network flows + a role about compliance, security, forensics etc.
- Failover routing by Health! only in the event of an outage.
- Weighted routing routes by a proportion you set
  ↳ to % to this other version. Canary deployment, A/B testing, blue/green.
- Multivalued answer
  ↳ returns up to 8 healthy records, health-checked.
  ↳ Search indexing or load balancer, step healthy and check one!

CVE = Common Vulnerabilities and Exposures
↳ Amazon Inspector scans for them.
↳ EC2 instances, container images
↳ EC2 and Container Registry

Cloudtrail is recently deprecated. Trail Insights is more operational learning

Block 3. Be introduced in role. CRL values can be modified live! you do not change the policy change. IAM and through ~~~~ with no downtime
↳ Role only ever increases. only the system does not change.

Day 10. A Sustained Effort.

Amplify

- a development service in AWS
- allows building of extensible, full-stack web and mobile apps.
- allows ML integration

Modules.

- Studio
- Libraries
- CLI
- Hosting

Serverless Computing

- on-demand service
- less server management
- FaaS: Lambda, Fargate, Aurora, DynamoDB, S3
  - function as a service
- Traditional; IaaS (Infrastructure as a Service; FaaS
- Lambda: max 15 minutes straight runtime.
  - start up a few ms to a second.
- powered by microVMs
- edge computing

Architectures

- Static Single Page Application
  - S3 + Cloudfront
- Containerised Application
  - Fargate
- Service-Oriented Architecture
  - Lambda + Gateway
- Serverless Architecture
  - Lambda, Gateway, Fargate, DynamoDB
- Serverless **Database**
  - DynamoDB
  - Aurora Serverless
  - for applications that have sporadic or infrequent database usage patterns
    - costs way less than a regular server-based database

Route 53 Overview

- a global DNS service
- provides different Routing Policies
- allows registering of your own domain name.
- create health checks, route traffic flows, configure DNS resolvers

Alias Record, Non-Alias Record

- Alias Record
  - Route Traffic to selected AWS resources
  - works like a CNAME Record
  - Not visible to DNS resolvers
  - points to a specific AWS resource
- Non-Alias Record
  - allows you to specify the IP addresses or the custom domain names of your servers or resources
  - visible to DNS resolvers
  - points to a particular IP address

Bind

- an open-source program
- internal DNS service

Active (Live) vs Passive (Failover)

- active-active improves fault tolerance and performance of your applications
- additional cost
- several active environments
- active-passive:
  - basic fault tolerance
  - cost-effective
  - one active environment

CloudFront Overview

- CloudFront -> CDN

Edge Location

- refers to the 'edge' or the boundary of the network.
- connects the different networks of various ISPs or Telecom companies.

Origin

- S3, ELB
- HTTP only
- HTTPS only
- Match Viewer

Distribution

Viewer

- HTTPS only
- Redirect HTTP to HTTPS
- HTTP and HTTPS

Security Features

- CDN
- Static, Dynamic
- Protocol Policy

OAI

- distributions with an S3 bucket as the origin
- you can associate to the Origin or Origin Group
- S3 bucket policy must be configured too.

Field-Level Encryption

- encrypt the specific data fields
- public key
- private key.

Sign URLs and Signed Cookies

- distributing private content
- signed url: https://tutorialsdojo.com/report.pdf ?Expires=1357034400 5 &Signature=nitfHRCrtzivO2HwPtWw~yYDhUF5EwRunQA... &Key-Pair-Id=K2JCJMDEHXQW5F
- signed cookies: in inspect.

Geo-Restriction

- 'allow-list'

Alternate Domain Name & SSL Certificate

- SNI = Server Name Indication
- Dedicated IP ADdress

CloudFront -> AWS

- WAF - CloudFront Integration
- AWS Shield

Comparisons

- High Availability vs Fault Tolerance.
- RTO vs RPO
- NACL vs Security Group
- File Gateway vs Volume Gateway vs Tape Gateway
- Integration vs Migration
  - Storage Gateway vs. DataSync
- EBS vs EFS vs S3
  - Durability
  - Access Method
  - Scalability
  - Latency
  - Backups
  - Data Encryption
  - Access Control
  - NFSv4 Protocol Support
  - Data LifeCycle
  - Use Cases

λambda-Testing

- cached volumes vs stored volumes
  - cached volumes keep only the hot subset on-prem with S3 as primary storage
    - stem demanded the *entire* dataset local
- AWS CloudTrail management events vs **data events**
  - Object-level calls are **data plane**; management events never record DeleteObject
- if the instance can deal with FAILOVERS, then it is HIGHLY AVAILABLE
  - if it has multiple clones, it is FAULT TOLERANT
    - Fault Tolerant => uptime of 100%
  - High Available means uptime of 99.99%
- CloudWatch can't tell you anything about Security-group changes.
- A lifecycle policy can move files not accessed in EFS to EFS Infrequent Access after 30 days.

Day 11, Sunday 08/09/24

Mock 3 | 49%. Time taken 70 minutes/ 180

Q1) MQ is for public events?...

Q2) AZ ⊆ Region
You can put auto-scaling between AZ!

Q3) ECS Cluster is not serverless
↳ choose Lambda instead

Q6) cloudformation template

Q7) I guess Kinesis Data Streams can run out time
– Data Exchange is just a marketplace

Q8) Polly can't do transcribing!
↳ it does text to low speech!

Q9) DynamoDB good for changing schemas

Q10) Active-Passive is when you don't need constant availability

Q12) AWS Shield is best for protecting against DDOS Attacks.
SSE = server-side encryption
Server-Side: SSE on S3 (AES-256)
Client-Side: encrypt at rest

Q13) S3 object is owned by the account that created it

Q17) ECRs are primarily used in EBS, not in Auto Scaling.
Target group, is wrong. "Launch Template"

Q18) NACL is more for controlling traffic in and out of the whole VPC.
"Still being able to reach other instances" is the key.

Q20) Use Reader endpoint for load-balanced support on read-only connections to the DB cluster. incl queries
↳ reduces overload on primary cluster

Q22) Be careful with fulfilling the information requirements of the Q. Hence 'cw' was wrong.

Q23) Apparently you can use DMS to convert server to serverless Aurora

Q24) Apparently DeleteOnTermination can be set to false on EC2. wtf?
↳ Ok! it's about the Volume.

Q25) Bastion hosts are always locked down. /0 means lots of IPs, /32 is a specific one

Q26) File Gateway presents a file-based interface to S3 which appears as a network file share. Allows NFS/SMB mount points ← file shares

Q27) To prevent DDoS, you can do:
– use cloudfront for distributing static/dynamic content
– ALB with Auto Scaling groups for EC2. Relevant alarms Inbound traffic to RDS by assigning it to a new private subnet.
– Set up alerts in cloudwatch to look for high Network In and CPU utilization metrics

Q28) Cluster placement group does allow low latency
SQS decoupling was the answer

Q30) for many domain names looked into one certificate you're going to have to upload all SSL certs of the domains in the ALB with the console and bind multiple certificates to the same secure listener on load balancer. Then ALB will automatically choose the optimal TLS cert for each client using SNI - Server Name Indication

Q32) By default, Auto scaling is not enabled in a DynamoDB table

Q34) You wouldn't be billed if your On-Demand instance is in a pending state. You would be billed though, if it was in a stopping state (preparing to hibernate)

Q37) Cluster - packs instances close together inside an availability zone. Low latency
Partition - Spreads instances across logical partitions such that groups of instances in one partition do not share the underlying hardware with groups of instances in different partitions.
Spread - Strictly places a small group of instances across distinct underlying hardware to reduce correlated failures

Q38, 41, 42, 43, 44, 46, 47, 49, 50, 52, 53, 54, 55, 56, 57, 59, 62, 63

Day 1a, Monday 02/09/24

Q39) 41, 42, 43, 44, 45, 46, 47, 48, 49, 50, 52, 53, 54, 55, 56, 57, 58, 59, 62, 63

Q39) Reserved Instances and On-Demand Capacity Instances are better suited for workloads with steady, predictable usage.

Q42) IAM User group is a group of IAM users.

- SCP cannot be attached to a role
- You also cannot associate an IAM role with an IAM user

Q42) Amazon Transit Gateway is for connecting VPCs and on-prem networks through a central hub.

- S3 Interface endpoint will be expensive
  ~ All the question numbers skipped!

Q1) You can't get KMS to automatically send a snapshot to S3.

Q8) Auto-scaling not auto-enabled in DynamoDB

Q3) ECS, not serverless

Q4) If you are already connect to the choice then your route table is fine

Q9) AppSync providing easyship to DynamoDB to keep easyship up-to-date

Q11) You gotta know or the field:

- An EC2 instance can communicate with the internet via a NAT instance/gateway. So it can also use an elastic IP.
- The allowed blocksize in VPC is 16/16 reserved and /27 ↑ 28.56 ✗ 16 RESERVED. It's actually 16 ≤ ≤ 28

Q16) DMS will have no downtime

Q17) DynamoDB Streams does 24hrs

Q23) NAT gateway can't attach throttling of a Lambda function

- ENI = Elastic Network Interface are virtual network interfaces that can be attached to instances or Lambda functions within your VPC.
  ↳ Lambda creates and manages ENI's when you run a Lambda function access to resources within your VPC.
  ↳ Therefore a Service Quota can be hit, you'll run out of IPs/ENI

Q23) EIP = Elastic IP.

- ENI shouldn't stay attached even if EC2 is stopped
- EIP should also remain associated
- Underlying host may have changed
  ↳ ENI's and data may be lost

Q27) Using KMS key must have a key policy
  ↳ Using key to control access to KMS keys one-to-one.
- You can also use IAM policies to control the key
  ↳ must be explicitly allowed by the key policy itself.
- Resource policy's specifies who can invoke the Lambda function, not which this specifies it can use.
  ↳ developer permission must then be added to the function's execution role, not on the resource policy.

Q33) SCP (Service Control Policies) are critical control over the maximum available permissions for all accounts in your organization.
  ↳ not sufficient, but necessary. (no permissions granted)
  ↳ DENY to block root access

Q39) OAC (Origin Access Control) on S3
- cloudfront signed URLs or signed cookies

Q36) Job bookmarking on AWS Glue is a legit feature. Reading data is a good idea in an ETL pipeline, in case we need to rerun the pipeline.

Q37) A VPN connection is not regardless of providing consistent and dedicated access to the on-prem network, having traffic the public internet and doesn't use a dedicated connection

Q44) Enable AutoScaling storage setting in MySQL RDS. solves the storage problem without further intervention

Q46) Use Multipart upload past 100MB in S3.

Q47) Provisioned Share is integrated with Secrets Manager

Q48) Access logging is an optional feature of ALB (Application load balancing)
  ↳ disabled by default.

Athena is serverless and can query easyship in S3.

Q52) It is not possible to enable/disable hibernation for an instance after it has been launched.

Q56) Aurora Read Replicas allow you to fully access data when occasional urgent sys or data was required.
  ↳ Provisioned capacity ensure that the read/write capacity is available when are need it.

Topic Test. Auto Scaling

- Target tracking = you name ONE number and nothing else. Pick 50% CPU. AWS authors the alarms, picks the adjustments, does the rest. The thermostat.
- Step scaling = you author the alarms AND a ladder of adjustments that vary with the size of the breach. If the storm hands you thresholds, alarms, or more than one adjustment, it is step.

Topic Test. DynamoDB

- you still cannot integrate DynamoDB table with CloudFront as these two are incompatible.
- A global secondary index can only be created in the region where its parent table resides.
- the more distinct partition key values your workload accesses, the more those requests will be spread across the partitioned space
- Application Load Balancer is not suitable to be used with DynamoDB
- you usually put EC2 instances on an Auto Scaling Group, and not a DynamoDB table.
- using a Network Access Control List alone is not enough to prevent traffic traversing to the public Internet. Moreover, you cannot copy DynamoDB on-demand backups to a different account or Region.
- you have to manually enable DynamoDB Streams first before you can use its endpoint
  - hence, "Amazon Kinesis Client Library (KCL) write an application that leverages on DynamoDB Streams Kinesis Adapter that will fetch data from the DynamoDB Streams endpoint. When there are updates made by a particular user, notify the subscribers via email using Amazon SNS " is incorrect.
    - it didn't first enable DynamoDB streams
  - same with this: "Create an AWS Lambda function that uses DynamoDB Streams Amazon Kinesis Adapter which will fetch data from the DynamoDB Streams endpoint. Set up an Amazon SNS Topic that will notify the subscribers via email when there is an update made by a particular user".
    - gotta have the Answer stem say that you ENABLE DYNAMODB STREAMS FIRST.

Topic Test: S3

- Site-to-Site VPN is just used for establishing secure connections between an on-premises network and an Amazon VPC

Topic Test: EFS

- using Bursting Throughput mode won't be able to sustain the constant demand of the global application
- EBS volumes are block storage devices attached to a single instance, making them unsuitable for shared storage.

Topic Test: EBS

- EBS volumes can only be attached to an EC2 instance in the same Availability Zone.
- AWS Backup can manage retention rules, but is typically used for centralized backup across multiple AWS services. It is not just focused on automating EBS snapshot creation. While possible, it is not the most cost-effective or straightforward solution for this specific use case.
  - Instead you should use "DLM (Data LifeCycle Manager)"
- Amazon S3 Server-Side Encryption and S3 Client-Side Encryption is incorrect because these relate only to S3.
  - on EBS, they use KMS bring-your-own OR Amazon managed AWS KMS.
- Instance store attached to EC2 Instance VS. EBS Volume attached to EC2 instance.
  - **the use of instance store volume is not suitable for mission-critical workloads because the data can be lost if the underlying disk drive fails, the instance stops, or if the instance is terminated.**
- an Instance Store only provides ephemeral storage where the data is automatically deleted when the instance is terminated
- Magnetic volumes provide the lowest cost per gigabyte of all EBS volume types and are ideal for workloads where data is accessed infrequently, and applications where the lowest storage cost is important.
- multi-attach feature can only be enabled on EBS Provisioned IOPS io2 block express or io1 volumes. In addition, multi-attach won't offer multi-az resiliency because this feature only allows an EBS volume to be attached on multiple instances within an availability zone.

Day 13.

Anki, done.

Section test: Design High-Performing Architectures

- Q1. By default, Lambda functions come with 512 MB of temporary storage, but you can configure up to 10 GB of storage
  - Amazon EFS for persistent storage
  - EC2 INSTANCES ARE NOT SERVERLESS!!!!
- Q7. difference between ALB and NLB?
  - UDP is not supported in ALB
  - UDP is Layer 4, i.e. network.
  - yes, DynamoDB for Key-value pair store.
- Q8. ECS is mainly used to host Docker applications
- Q9. Recall that **Rekognition** is for image recognition.
- Q11. Gotta deploy the bastion host in a Public subnet of a VPC
- Q12. AWS CloudFormation provides a common language for you to describe and provision all the infrastructure resources in your cloud environment.
- Q13. EBS has a storage option called "Cold HDD," but due to its price, it is not ideal for data archiving. EBS Cold HDD is much more expensive than Amazon S3 Glacier / Glacier Deep Archive and is often utilized in applications where sequential cold data is read less frequently.
- Q14. You can use the consolidated billing feature in AWS Organizations to consolidate payment for multiple AWS accounts or multiple AISPL accounts. With consolidated billing, you can see a combined view of AWS charges incurred by all of your accounts. You can also get a cost report for each member account that is associated with your master account.
  - For this particular scenario, it is suitable to use AWS Organizations instead of setting up an AWS Transit Gateway since the objective is for maintaining administrative control of the AWS resources and not for network connectivity.
- Q16. LambdaEdge lets you run Lambda functions to customize the content that CloudFront delivers, executing the functions in AWS locations closer to the viewer.
  - origin failover by creating an origin group with two origins with one as the primary origin and the other as the second origin which CloudFront automatically switches to when the primary origin fails. This will alleviate the occasional HTTP 504 errors that users are experiencing.
    - although establishing multiple VPCs in different AWS regions and configuring a transit VPC to interconnect them IS VIABLE. **It is not the most cost-effective option.** Use **LambdaEdge** instead!
- Q17. AWS Batch jobs does not provide the interactive, high-performance SQL analytics or BI integration needed for the semi-structured and structured data.
  - it's only compute.
- Q18. SQS is a message queuing service that ensures message delivery between distributed components, but it's not optimized for real-time processing of high-velocity data streams. It's better suited for tasks that can handle a bit of latency.
- Q19. A Gateway Load Balancer **does not support path-based routing**
- Q20. Amazon FSx still has a higher level of throughput and IOPS compared with AWS Storage Gateway
- Q21. Amazon FSx for NetApp ONTAP is a fully managed AWS service that provides high-performance, scalable file storage based on NetApp's ONTAP file system. It offers versatile storage options, supporting both file (NFS, SMB) and block (iSCSI) protocols, making it compatible with Windows, Linux, and macOS environments.
  - Amazon S3 is primarily an object storage, not block storage. It doesn't offer the low-latency access required for a trading application.
- Q24. AWS VPN CloudHub is only for VPNs and not for VPCs.
  - It is also not capable of managing hundreds of VPCs with multiple VPN connections to their data centers that span multiple AWS Regions.
- Q25. Storage Optimised Instances!
  - Memory Optimised -> in-memory
  - Compute Optimised -> compute bound. batch processing, media transcoding
  - General Purpose -> too basic.
- Q29. you cannot only set a priority to individual items in the SQS queue!
- Q30. encryption at rest does not affect the traffic routing
  - AWS Direct Connect is primarily used to establish a dedicated network connection from your premises to AWS
  - AWS VPN CloudHub is typically used to provide secure communication between remote sites and not for creating a private endpoint to access Amazon S3 and DynamoDB within the Amazon network.

Topic-test: ELB (elastic load balancing)

- AWS Firewall Manager just simplifies your AWS WAF and AWS Shield Advanced administration and maintenance tasks across multiple accounts and resources.
- Although NACLs can help you block incoming traffic, this option wouldn't be able to limit the number of requests from a single IP address that is dynamically changing.
- Since the security group is properly configured, the issue may be caused by a wrong health check configuration in the Target Group.
- **Application Elastic Load Balancer**: only allows 2 types of health check: HTTP and HTTPS.
  - A TCP health check is only offered in Network Load Balancers and Classic Load Balancers.
  - ICMP health check and FTP health check are incorrect as these are not supported.
- AWS CloudTrail is primarily used to monitor and record the account activity across your AWS **resources** and not your web applications
- Network Load Balancers do not support gRPC
- When cross-zone load balancing is enabled, each load balancer node distributes traffic across the registered targets in all enabled Availability Zones

Topic Test: IAM (Identity & Access Management)

- IAM Identity Center is the recommended approach for workforce authentication and authorization on AWS for organizations of any size and type, but not for generating tokens.
- When you use the AWS Management Console to create a user, you must choose to include at least a console password or access keys. By default, **a brand new IAM user created using the AWS CLI or AWS API has no credentials of any kind.**
- If your identity store is not compatible with SAML 2.0 then you can build a custom identity broker application to perform a similar function.
  - simply using an IAM policy is not enough to integrate your LDAP service to IAM. You need to use SAML, STS, or a custom identity broker.
  - IAM Identity Center uses SAML 2.0

Topic Test: Lambda

- overkill to use AWS Glue for processing 1 MB of data per day. AWS Glue has a minimum billing duration of 1 minute (Glue version 2.0 and later), which is not cost-effective for processing small amounts of data quickly (in 10 seconds).
- Data Firehose is meant for delivering streaming data to a storage
- 1 MB per day), using Redshift would be a poor fit.
- With AWS Lambda, you pay only for what you use.
  - unlike EC2 which is running all the time.
- Amazon Athena is just an interactive query service that enables you to easily analyze data in Amazon S3 using standard SQL. Kinesis should be used to process the streaming data in real time.

Review Mock 4

- Q1. The Fleet Manager is just one of the many capabilities of the AWS Systems Manager that provides a unified user interface (UI) experience in remotely managing your nodes/EC2 instances running on AWS or on-premises.
  - AWS RAM is used to securely share your resources across AWS accounts in your AWS organization and not for tracking or controlling software licenses.
- Q4
  - ```
    {
      "Version":"2012-10-17",
      "Statement":[
        {
          "Sid":"DirectoryTutorialsDojo1234",
          "Effect":"Allow",
          "Action":[
            "ds:*"
          ],
          "Resource":"arn:awscdsus-east-1:987654321012:directoryds-1234567890"
        }
      ]
    }
    ```
    Statement ID
    Account ID
    Directory ID
- Q5.
  - When you launch an EC2 instance into a default VPC, AWS provides it with public and private DNS hostnames that correspond to the public IPv4 and private IPv4 addresses for the instance.
    - However, when you launch an instance into a non-default VPC, AWS provides the instance with a private DNS hostname only. New instances will only be provided with a public DNS hostname depending on these two DNS attributes: the DNS resolution and DNS hostnames that you have specified for your VPC and if your instance has a public IPv4 address.
    - In this case, the new EC2 instance does not automatically get a DNS hostname because the DNS resolution and DNS hostnames attributes are disabled in the newly created VPC.
- Q6.
  - AWS Application Migration Service (AWS MGN) is the primary migration service recommended for lift-and-shift migrations to AWS.
  - AWS Application Discovery Service is primarily used to track the migration status of your on-premises applications from the Migration Hub console in your home Region. This service is not capable of doing the actual migration.
- Q7. Standby replica is used in Multi-AZ deployments
- Q8. The IAM certificate store is an *older feature* that can also store SSL/TLS certificates.
- Q11. While Amazon RDS Multi-AZ deployments provide high availability and automatic failover, they primarily focus on database redundancy rather than improving website performance.
- Q12. testing on a local machine may not accurately replicate the resources and environment of a large RDS DB instance. Thus, the tests might produce unreliable results.
  - note that a stopped DB instance does continue to incur costs on its provisioned storage.
- Q14. Even though S3-IA costs less than the S3 Standard storage class, it is still more expensive than S3-One Zone IA
- Q16. Indeed no cost for a stopped EC2 instance
  - but there is for a *running* EC2 instance as well as for an EBS volume attached to a stopped EC2 instance.
- Q19.
  - io1 = 4GB - 16TB
  - 100 IOPS - 64,000 IPOS
  - **ratio of 50:1**
- Q22. Although Amazon Redshift is a fully-managed data warehouse, you will still need to configure cross-region snapshot copy to ensure that your data is properly replicated to another region.
- Q23. Predictive scaling uses machine learning to predict capacity requirements based on historical data from CloudWatch.
- Q25. EBS is primarily used as a storage of your EC2 instances.
  - silly mistake!
  - storage gateway: "on-prem"
- Q27. you cannot directly use a CloudFront distribution ID as a principal in an S3 bucket policy
  - OAC = Origin Access Control
- Q29. An Amazon S3 Glacier (Glacier) vault can have one resource-based vault access policy and one Vault Lock policy attached to it.
  - enforce regulatory and compliance requirements.
- Q31. you can configure a bucket to be a Requester Pays bucket
  - note that if you enable Requester Pays on a bucket, anonymous access to that bucket will not be allowed anymore.
- Q36. Amazon EKS supports two autoscaling solutions:
  - Karpenter is a flexible, high-performance Kubernetes cluster autoscaler that launches appropriately sized compute resources, like Amazon EC2 instances
  - Cluster Autoscaler, on the other hand, automatically adjusts the number of nodes in a cluster based on pod failures or rescheduling
- Q37. A global secondary index can only be created in the region where its parent table resides
- Q38. An Elastic Fabric Adapter (EFA) is a network device that you can attach to your Amazon EC2 instance to accelerate High Performance Computing (HPC) and machine learning applications
  - Elastic Network Adapter (ENA) doesn't have OS-bypass capabilities, unlike EFA.
  - Elastic Network Interface (ENI) is simply a logical networking component in a VPC that represents a virtual network card. It doesn't have OS-bypass capabilities
- Q39. you can't create an alarm in Amazon SNS.
- Q40. it is not possible to install the software in Amazon S3.
- Q45. custom Config rule is just unnecessary since a managed rule already exists for access key rotation
  - Amazon EventBridge cannot directly check for IAM events that contain the age of IAM access keys. You must use the AWS Config service to detect the non-compliant IAM keys.
- Q46. AWS Service Catalog is not used to detect non-compliant resources and is only used to manage catalogs of IT services from virtual machine images, servers, software, databases, and other resources.
- Q48. a standard queue only provides at-least-once delivery and does not guarantee strict message ordering.
  - Kinesis data stream is composed of shards, where each shard stores a sequence of ordered data records
- Q50. Amazon S3 Glacier is used for data archive only.
- Q51. An IAM Role, in itself, cannot directly make the S3 objects public or change the permissions of each individual object.
  - CORS will only allow objects from one domain (travel.celebu.com) to be loaded and accessible to a different domain (palawan.com). It won't necessarily expose objects for public access all over the internet.
- Q53. the requirement is to send notifications for upcoming or scheduled AWS events. This solution will send notifications after the EC2 instances have stopped and not show upcoming events in AWS.
- Q54. AWS AppConfig is primarily designed to manage and deploy application configuration changes in a controlled and safe manner. It does not address the key challenges of schema conversion, data migration, or SQL query compatibility crucial for transitioning from Microsoft SQL Server to Amazon Aurora PostgreSQL.
  - **where an organization is transitioning from Microsoft SQL Server to Amazon Aurora PostgreSQL, Babelfish is a crucial tool**
- Q58. an instances IP address is Metadata, not user data.
  - http://169.254.169.254/latest/meta-data/
- Q62. The PII (Personally Identifiable Information) redaction feature of the Amazon Textract Medical service is not enough to identify all the Protected Health Information (PHI) in the PDF reports.

ARN = Amazon Resource Name
    ```

Day 14. Wednesday.

Topic Test: RDS

- Q2. AWS Schema Conversion Tool is typically used for heterogeneous migrations where you're moving from one type of database to another (e.g., Oracle to PostgreSQL)
- Q5. although you can create and assign an IAM Role to your EC2 instances, you still need to configure your RDS to use IAM DB Authentication.
  - IAM database authentication works with MySQL and PostgreSQL. With this authentication method, you don't need to use a password when you connect to a DB instance. Instead, you use an authentication token.
- Q6. transparent data encryption (TDE) is primarily used to encrypt stored data on your DB instances running Microsoft SQL Server and not the data that is in transit.
  - IAM database authentication is only supported in MySQL and PostgreSQL database engines
- Q7. Is RDS ACID?
  - Standby replica is used in Multi-AZ deployments

Section Test: Design Secure Architectures

- In order for you to establish an SSH connection from your home computer to your EC2 instance, you need to do the following:
  - On the Security Group, add an Inbound Rule to allow SSH traffic to your EC2 instance.
  - On the NACL, add both an Inbound and Outbound Rule to allow SSH traffic to your EC2 instance.
- Q4. Amazon Workspace to create the needed virtual desktops in your VPC.
- Q5. Amazon EBS does not support object lock
  - EFS only supports file locking

EBS vs. EBS Volume? same thing?

- Q8. what does ElastiCache do?
  - Using Redis AUTH command can improve data security by requiring the user to enter a password before they are granted permission to execute Redis commands on a password-protected Redis server.
    - **IAM authentication is simply not supported for executing Redis commands like MULTI EXEC, and IAM tokens expire every 12 hours, which does not align with the need for long-lived credentials.**
  - Redis vs other DBs? is it relational? vs Aurora? vs RDS?

Is Neptune a Graph database?

- Q13. Origin Shield feature is not primarily used for security but for improving your origin's load times
  - Require that your users access your Amazon S3 content by using CloudFront URLs, not Amazon S3 URLs.
- Q14. SAML 2.0-Based Federation by using a Web Identity Federation is incorrect because this is primarily used to let users sign in via a well-known external identity provider (IdP), such as Login with Amazon, Facebook, or Google. It does not utilize Active Directory.
- Q15. A Gateway endpoint is a type of VPC endpoint that provides reliable connectivity to Amazon S3 and DynamoDB without requiring an internet gateway or a NAT device for your VPC.
  - When you create a Gateway endpoint, you can attach an endpoint policy that controls access to the service to which you are connecting
  - It would take you **a lot of time** to configure a bucket policy for each S3 bucket instead of using a single endpoint policy.
- Q16. bucket policies and access control lists (ACLs) are resource-based policies
  - You can also manage the public permissions of your objects during upload.
  - CORS will only allow objects from one domain (travel.cebu.com) to be loaded and accessible to a different domain (palawan.com). It won't necessarily expose objects for public access all over the internet.
  - CORS will only allow objects from one domain (travel.cebu.com) to be loaded and accessible to a different domain (palawan.com). It won't necessarily expose objects for public access all over the internet.
- Q18. EBS Encryption!
  - Snapshots of an **encrypted volume** are automatically encrypted.
  - All data moving between the volume and the instance are encrypted.
- Q20. you can't block IP Addresses from a Security Group. I think you can do so from **Network ACcess Control List though**
- Q22. A key policy is a resource policy for an AWS KMS key
  - Key policies are the primary way to control access to KMS keys
  - All Amazon FSx for NetApp ONTAP file systems is encrypted at rest with keys managed using AWS Key Management Service (AWS KMS).
  - **the resource policy specifies who can invoke the Lambda function, not which AWS operations it can use.**
- Q23. store the responses in an Amazon Kinesis data stream and set its retention period
  - Create an Amazon Pinpoint journey for the multi-engagement SMS marketing campaign and an Amazon Kinesis Data Stream for analysis. Configure Amazon Pinpoint to send events to the Kinesis data stream for collection, processing, and analysis. Set the retention period of the Kinesis data stream to 365 days.
- Q25. Port 3389 is for **RDP!**
  - INGRESS vs. EGRESS

Recall that:
AWS Dataprep is primarily used to simplify your migration with AWS, it makes it simple and fast to move large amounts of data between on-premises storage and Amazon S3 or Amazon Elastic File System (Amazon EFS).

- Q27. EBS volumes are for block storage, and while they primarily focus on disk-level encryption, S3 requires object-level encryption
- Q28. A VPC endpoint enables you to privately connect your VPC to supported AWS services and VPC endpoint services powered by PrivateLink without requiring an Internet gateway, NAT device, VPN connection, or AWS Direct Connect connection.
  - Endpoints are virtual devices. They are horizontally scaled, redundant, and highly available VPC components that allow communication between instances in your VPC and services without imposing availability risks or bandwidth constraints on your network traffic.
  - Internet gateway is used for instances in the public subnet to have accessibility to the Internet.
  - NAT Gateway allows instances in the private subnet to gain access to the Internet, but not vice versa.
- Q29. Match Viewer is an Origin Protocol Policy that configures CloudFront to communicate with your origin using HTTP or HTTPS, depending on the protocol of the viewer request.
- **Q30. Remember that you won't get billed if you use a Gateway VPC endpoint for your Amazon S3 bucket, unlike an Interface VPC endpoint that is billed for hourly usage and data processing charges. Take note that the scenario explicitly asks for the most cost-effective solution.**

Topic Test: SQS

- Q1. Because Amazon SQS is a distributed system, there's no guarantee that the consumer actually receives the message
  - Thus, the consumer must delete the message from the queue after receiving and processing it.
- Q2. Amazon SQS uses short polling by default, querying only a subset of the servers (based on a weighted random distribution) to determine whether any messages are available for inclusion in the response. Short polling works for scenarios that require higher throughput. However, you can also configure the queue to use Long polling instead, to reduce cost.
- Q3. Amazon SQS automatically deletes messages that have been in a queue for more than the maximum message retention period. The default message retention period is 4 days.
  - you can increase the message retention period to a maximum of 14 days using the SetQueueAttributes action.
- Q4. you cannot only set a priority to individual items in the SQS queue
- Q5. fan-out the messages to multiple SQS queues using a filter policy in Amazon SNS subscriptions to allow parallel asynchronous processing. By doing so, the entire message processing will not exceed 24 hours.
- Q6. The visibility timeout is a period of time during which Amazon SQS prevents other consuming components from receiving and processing a message.
  - immediately after the message is received, it remains in the queue.
    - To prevent other consumers from processing the message again, Amazon SQS sets a visibility timeout, a period of time during which Amazon SQS prevents other consumers from receiving and processing the message. The default visibility timeout for a message is 30 seconds. The maximum is 12 hours.
  - **Component Timeout** is not a concept or configurable parameter that exists in Amazon SQS
  - **Receive message wait time**, which controls short or long polling behavior; it has nothing to do with preventing other consumers from processing the same message.
- Q7. Simple Workflow Service = **SWF**
  - Amazon Simple Queue Service (SQS) and Amazon Simple Workflow Service (SWF) are the services that you can use for creating a decoupled architecture in AWS.
  - you can't create a VPC peering for your on-premises network and AWS VPC.
- Q8. Amazon S3 notification feature enables you to receive notifications when certain events happen in your bucket
  - Amazon S3 supports the following destinations where it can publish events:
    - Amazon Simple Notification Service (Amazon SNS) topic
    - Amazon Simple Queue Service (Amazon SQS) queue
    - AWS Lambda
  - you can only add 1 SQS or SNS at a time for Amazon S3 events notification. If you need to send the events to multiple subscribers, you should implement a message fanout pattern with Amazon SNS and Amazon SQS.
  - **you can't poll Amazon SNS**. Instead of configuring queues to poll Amazon SNS, you should configure Amazon SQS queue to subscribe to the SNS topic.

Topic Test: CloudFront

- Q1. S3 presigned URLs provide temporary access to specific files
- Q2. CloudFront is primarily a content delivery network service that speeds up the delivery of content to your customers.
  - WorkDocs is simply a fully managed, secure content creation, storage, and collaboration service.
- Q4. You can control how long your objects stay in a CloudFront cache before CloudFront forwards another request to your origin. Reducing the duration allows you to serve dynamic content.
- Q6. Use Amazon ElastiCache for the website's in-memory data store or cache.
  - I had no idea you could use ElastiCache for website speed. I thought it was only for DB?
- Q7. Global Accelerator is incorrect because this service is more suitable for non-HTTP use cases
- Q8. WAF for DDoS! Not GuardDuty

Section Test: Cost Optimised

- Q3. AWS Elastic Disaster Recovery (AWS DRS) provides continuous block-level replication, recovery orchestration, and automated server conversion capabilities. These allow customers to achieve a crash-consistent recovery point objective (RPO) of seconds, and a recovery time objective (RTO) typically ranging between 5–20 minutes.
- Q4. S3 Object Lock in compliance mode, which is essential to prevent deletion or modification during the retention period.
  - governance mode is incorrect because it does not only provide the same level of protection as compliance mode
- Q9. governance mode is incorrect because it does not only provide the same level of protection as compliance mode
- Q10. Data Firehose is a managed delivery service optimized for buffering and loading data to destinations (like S3)
- Q14. only Amazon S3 Standard has the feature of no minimum storage duration. It is also the most cost-effective storage service because you will only be charged for the last 12 hours, unlike in other storage classes where you will still be charged based on its respective storage duration (e.g. 30 days, 90 days, 180 days). S3 Intelligent-Tiering also has no minimum storage duration and this is designed for data with changing or unknown access patterns
  - Data that is deleted from S3 Standard-IA within 30 days will still be charged for a full 30 days.
    - basically the same for Glacier Deep Archive
- Q15.
  - You will be billed when your On-Demand instance is preparing to hibernate with a *stopping* state
  - You will be billed when your Spot instance is preparing to stop with a *stopping* state
- Q18. IAM policies are primarily used to specify what actions are allowed or denied on your S3 buckets. You cannot configure an IAM policy to automatically purge logs for you in any way.
- Q19. Storage Gateway is mainly used in providing low-latency access to data by caching frequently accessed data on-premises while storing archive data securely and durably in Amazon cloud storage services.
  - Storage Gateway optimizes data transfer to AWS by sending only changed data and compressing data.
  - only Amazon S3 Standard has the feature of no minimum storage duration. It is also the most cost-effective storage service because you will only be charged for the last 12 hours, unlike in other storage classes where you will still be charged based on its respective storage duration (e.g. 30 days, 90 days, 180 days). S3 Intelligent-Tiering also has no minimum storage duration and this is designed for data with changing or unknown access patters.
  - you didn't need continuous access to the Archived data.
- Q21. There are separate permissions for the use of a KMS key that provides added protection against unauthorized access of your objects in Amazon S3. SSE-KMS also provides you with an audit trail that shows when your KMS key was used and by whom.
- Q26. Even though S3-IA costs less than the S3 Standard storage class, it is still more expensive than S3-One Zone-IA.
- Q27. S3 Glacier Flexible Retrieval is typically more expensive than directly backing it up to Glacier Deep Archive
- Q29. Target Tracking is **dynamic**

Day 15

Schedule
- 1 Review
- 1 mock
- 1 cover that

* Review 42/65

- Q3, Hardware Security Module keys can be both permanently if you block it have a copy.

- Q5, Visibility Timeout is correct for the SQS queue

- Q6, is only-way outbound security group rules are implicit because of statefulness.
  ↳ the correct answer (obviously) was "Allow outbound on the NACL"

- Q7, Aurora. SQLite DB table
  ↳ Redshift = OLAP
  ↳ DynamoDB is ACID, but can't do OLTP & complex queries.
  ↳ RDS can't handle SQLite

- Q8) EFA doesn't work on windows. ENA does.
  "Parallel Cluster" is just a way to deploy and manage HPC clusters. It does not help with bandwidth.

Q11) AWS Config is more for compliance purposes as compared to real-time monitoring of quota usage.
  ↳ Trusted Advisor → has a Developer Support plan
    ↳ can publish service limits to CloudWatch
    ↳ can also use boolean functions to read data

Q13) Network Load Balancer doesn't have Weighted Target Groups.

Q14) Moving EC2 to private subnet makes the instances inaccessible to you

Q15) Run command. → used to configure EC2 without having to establish RDP or SSH to each instance
  ↓
  Applicaner provides a managed service for deploying and monitoring configurations into EC2 and API's from source code / container images.

Q17) Create an A record aliased to the load balancer DNS name.
  ↳ IP Address of the load balancer can change at any time.

Q18) CloudWatch is for systems monitoring and server metrics. CloudTrail records API calls and AWS management console actions.

Q19) App Sync provides resolvers solve the common problem of aggregating data from multiple datasources & tables.
  ↳ Athena available if allow write operations.

Q20) Karpenter is a flexible, high performance Kubernetes cluster autoscaler that launches appropriately sized compute resources (like EC2 instances)
  ↓ ↳ "Scale in and scale out" → refers to horizontal scaling
  EKS supports 2 autoscaling products:
  - Karpenter
  - Cluster Autoscaler

Q22) Any IPv4 Addresses would include the blocked IPv6 addresses! Need to be specific with these 192.220... is not explicitly allowed, so the ip will be blocked and allow!

Q23) The EBS snapshot can be used under the snapshot is still progress!
  ↳ not just read only

Q24) With an EBS volume (ENCRYPTED)
  1) snapshots are also encrypted
  2) All data transferred between the volume & instance is also encrypted.

Q30) Bulk retrieval on Glacier takes 5-12 hours.
  ↳ S3 One-zone IA costs more than Glacier

Q28) S3 already monitors an index of object key names in each AWS region. S3 stores key names in alphabetical order, the key name determines which partition the key is in.
  ↳ i.e. You don't have to do anything and the S3 backend will organize them.

Q29) Lifecycle rules. Pending → Pending:wait → Pending:proceed → Running → Terminating → Terminating:wait → Terminating:proceed → Terminated.

Q32) Network Access Analyzer is a feature of VPC that reports on unintended access to your AWS resources based on security and compliance.
  ↳ It does not capture of packets
  ↳ does provide information on traffic entering or leaving VPC
  ↳ AWS Network Firewall on the right! updates policy tables accordingly.

Q43) Gateway endpoint is the cheap one for S3/DynamoDB

Q44) Although Fargate is serverless Outposts is not!
  AppSync

Q47) Send all log events to Amazon Kinesis then develop a client processes to apply scrubbing on the logs.
  ↳ CWL doesn't have real-time processing capability

Q46) By default, CloudTrail event log files are encrypted using Amazon S3 SSE → server-side encryption
  ↳ Note they're stored in S3, not Glacier.

Q49) Equal cost Multipath routing

Q51) Create a WebACL rule in AWS WAF to block the specified country. Associate the rule to the ALB.
  ↳ using NACLs to block IPs is rather "block-a-mole"

Q52) Even tokens in a DB should go through IAM DB authentication
  ↳ RDS is always tricky with authentication services!

Q54) Aurora Replicas will flip the CNAME.
  ↳ Aurora itself will attempt to create a new DB instance in the same AZ as the original instance
  ↳ best-effort basis.

Q55) Set up VPC peering and re-configure the route table's target and destination of the instance's subnet.
  ↳ Note that VPC endpoints are region specific only and do not support inter-region communication

Q60) A scaling policy based on the CPUUtilization metric is not meant for time-sensitive messages since you need to ensure that the messages are processed within a ~~specific~~ specific time period
  ↳ Scale-out triggered on CPUUtilization of current instances, and not based on the age of the message