AWS Cheatsheet - Part 2

Author: Bhaskar S
Version: 1.8
Date: 03/02/2024
Blog: PolarSPARC

EC2

- EC2 on-demand fixed price by second, for short irregular uninterrupted workloads
- EC2 savings plan 1 or 3 yr, 72% off on-demand, price by hour, for consistent usage
- EC2 reserved for type/region/tenancy/platform 1 or 3 yr, lower hourly price, for predictable steady workloads
- EC2 spot 90% off on-demand, runs until terminated or price exceeds bid price, 2 min alert before termination, for flexible start/end work
- EC2 dedicated host for compliance/regulatory needs, legacy software licenses (per-socket, per-core, per-vm)
- EC2 dedicated instance share hardware with others in same account
- EC2 fleet and Spot fleet allow users to launch instances with mix of options (on-demand, savings, reserved, spot) in multiple AZs
- EC2 cluster placement in one AZ for low-latency, network performance workloads
- EC2 spread placement in different racks within an AZ OR different AZs in a region
- EC2 partition placement is logical segments in different AZs within a region
- EC2 Instance Store - temporary, ephemeral storage, for frequently accessed cache/scratch data, can be attached only at launch time **NOT** after
- EC2 AMI copied from one Region to a second automatically creates a Snapshot in the second Region
- EC2 scale-in policy will terminate instances launched with the oldest launch configuration first followed by oldest launch template next
- EC2 using a launch configuration has a default value null for the instance placement tenancy and controlled by the tenancy attribute of the VPC
- EC2 instance that is impaired can be automatically recovered using CloudWatch Alarms to be identical to the original instance, including the instance ID, private IP addresses, Elastic IP addresses, and all instance metadata

- Elastic IP incurs cost per hour if **NOT** attached, per region only, 5 max per account
- ENI is virtual NIC, each instance assigned a default primary and cannot be detached, can create and attach multiple secondary to an instance
- Security Group is Region specific, bound to ENI of an instance, like virtual firewall, stateful, only allow rules

EBS/EFS/FSx/Storage Gateway

- EBS (gp3, io1, io2), raw block device, bootable volume, in single AZ, attached to single instance over network

- EBS persists independent of instance, auto replicated within AZ, data intact after instance termination if 'Delete on Termination' option disabled
- gp3 (general purpose) for balance of price/perf for a wide range of transactional workloads, size 1 GB to 16 TB
- gp3 baseline IOPS 3000, maximum IOPS 16000, baseline throughput 125 MBps, max throughput 1000 MBps
- io1 (provisioned iops) for IOPS/throughput intensive workloads, size 4 GB to 16 TB, maximum IOPS 64000, supports up to 16 multi-attach within AZ
- io2 (provisioned iops block express) only for Nitro systems, size 4 GB to 64 TB, maximum IOPS 256000, supports up to 16 multi-attach within AZ
- st1 (throughput optimized HDD) not bootable, no multi-attach, size 125 GB to 16 TB, maximum throughput 500 MBps, for big data/data-warehouse workloads
- sc1 (cold HDD) not bootable, no multi-attach, size 125 GB to 16 TB, maximum throughput 250 MBps, lowest cost for less frequent access, data archiving
- EBS snapshot is point-in-time copy of EBS volume stored in S3, initial is full backup, subsequent is incremental backups
- EBS snapshot for creating identical copy in another AZ OR allows **COPY** to another region (with encryption)

- **Data Lifecycle Manager** for regular scheduled automatic creation/retention/deletion of EBS snapshot/AMIs, for DR as well as compliance

- EFS is a regional managed NFS, **ONLY** works in Linux, can attach to many EC2 instances in different AZ in same region
- EFS auto scale, pay per GB use, petabyte scale, allows encryption in-transit and at-rest, enable encryption **ONLY** at creation not after
- EFS Max I/O performance mode is used to scale to higher levels of aggregate throughput and operations per second
- EFS can be used for both AWS services as well as on-prem
- EFS Infrequent Access provides price/performance that is cost-optimized for files, not accessed every day, with cost up to 90% lower compared to EFS Standard

- FSx managed 3rd-party file systems like NetAPP, OpenZFS, Windows File Server, Lustre
- FSx integrates well with S3
- FSx Lustre for compute-intensive, high-performance, ML/HPC workloads, support for parallel and distributed processing
- FSx Windows File Server allows native Windows NTFS and SMB, with multi AZ support

- **Storage Gateway** connects on-prem storage with cloud storage for seamless/secure integration, for backup/restore, DR
- **S3 File Gateway** makes S3 buckets accessible via the NFS/SMB protocols for on-prem apps, caches most recently used files
- **FSx File Gateway** provides native access to FSx for Windows File Server for SMB clients on-prem, caches the most recently used files
- **Volume Gateway** provides iSCSI interface for S3 and allows for snapshots of on-prem volumes to the cloud

- **Tape Gateway** mimics the tape backup interface for S3 and allows the backup of on-prem data using leading vendor backup software
- Tape Gateway encrypts data between the gateway and AWS for secure data transfer, compresses data and transitions virtual tapes between S3, S3 Glacier, or S3 Glacier Deep Archive

ALB/NLB/GWLB

- ELB multiple targets in multiple AZ (select at least two), monitors health of targets, routes only to healthy targets
- ELB default security group allows http/s, provides ssl/tls termination, session stickiness

- ALB layer 7 (application), supports (http/2, http/s, grpc, websocket), url redirecting
- ALB supports routing (path, hostname, query string, headers), health checks, session affinity

- NLB layer 4 (transport), tcp, udp, tls, millions of request per sec, health checks, one static (or elastic) IP per AZ, health checks
- NLB will route traffic to instances using the primary private IP addr specified in the primary network interface for an instance if the targets are specified using an instance ID

- GWLB layer 3 (network), ip packets across all ports, health checks, integrates with network appliances using GENEVE protocol

- Cross Zone LB - for ALB option is **ALWAYS** enabled and no cost for data transfer, for NLB disabled and **WILL** incur cost for inter AZ data transfers

- Auto Scaling Group (ASG) min and max instances, launch template, health monitor to replace terminated or impaired instances
- ASG scaling policy - manual for simple workloads
- ASG scaling policy - scheduled according to predictable usage with desired, min, and max capacity on specific date/time

- Target Tracking Scaling - increase/decrease current capacity based on a monitoring metric and a target value
- Step Scaling - increase/decrease current capacity based on Alarm (Ex: if CPU > 70%, increase by 2 and if CPU < 30, decrease by 2)
- Simple Scaling - similar to Step Scaling, but with a Cool Down period (Default is 5 mins and no scaling during this period even if alarm)
- Predictive Scaling - uses ML to predict capacity based on historical usage, for situations that have cyclical pattern

- Sticky Session - for ALB **ONLY**, LB generated cookie name AWSALB with expiry of 7 days (not configurable), application generated cookie name must be registered with LB with expiry between 1 sec to 7 days

- Server Name Indication is an extension to TLS that allows clients to indicate the domain name they want to reach on the initial handshake, solves the problem of loading multiple SSL certificates onto

one web server (to associate multiple domains with a single IP address and TCP port), only works with ALB/NLB/CloudFront

VPC

- VPC isolated virtual private network, region scope, span AZs, non-overlapping ip4 cidr between /16 and /28, up to 5 max cidrs
- VPC cidr block size **CANNOT** change once created, can have one or more private/public subnets with IP addrs from the cidr block(s)
- VPC subnet has AZ scope (reside entirely in one AZ), subnet logical container for resources
- VPC has one **MAIN** route table (implicit vpc router) that handles routing within and outside vpc
- VPC can have custom route table, each subnet **MUST** be associated with **ONLY** one route table (MAIN or custom)
- Default VPC in a region has cidr **172.31.0.0/16**, with a public subnet in each AZ, with internet connectivity enabled
- Internet Gateway (IGW) scaled/redundant/available VPC component, attached to **ONLY** one VPC at a time, supports ip4 and ip6

- IGW connects resources in public subnet to internet AND vice-versa, entry in route table to IGW

- NAT Instance allows outgoing from private subnet to internet, prevents connections from internet, create/launch from public subnet
- NAT Instance must have elastic ip, source/destination option disabled, private subnet route table entry, customer managed, cannot auto scale
- NAT Instance can be used as a bastion server, supports port forwarding, can associate Security Groups

- NAT Gateway is managed service, has AZ scope, create/launch from public subnet, must have elastic ip, no security group
- NAT Gateway **CANNOT** be used as a bastion server, **DOES NOT** support port forwarding, **CANNOT** associate Security Groups

- Egress Only Internet Gateway similar to NAT Gateway but is for ip6

- Network Access Control List (NACL) is stateless firewall, works at subnet, one subnet one NACL
- Default NACL accepts in/out for the subnet associated with
- NACL rule numbers 1 to 32766, lower rule number higher priority, first rule match drives decision, allow/deny rules, for blocking specific ip addr to subnet

- **VPC Peering** for private route between two VPCs, ip4 and ip6, works for two VPCs in two regions, works in same or two different accounts
- VPC Peering is **NOT** transitive, for EC2 instances in two VPC subnets to communicate route table in both subnet must be updated
- VPC Peering **CANNOT** share IGW, NAT Instance, NAT gateway

- **VPC Endpoint (EP)** (PrivateLink) access to public services (s3, dynamodb, lambda) via AWS private network, two types - **Interface EP** and **Gateway EP**

- Interface EP needs ENI with private addr, uses DNS entries to redirect, must have security group, ONLY option for access to on-prem, cost per hour/per GB of data
- Gateway EP must be target in route table, uses prefix list to redirect, uses VPC EP policies, support for s3 or dynamodb **ONLY**, no cost

- **Site-to-site VPN** uses ipsec one duplex tunnel of 1.25 Gbps (outgoing, incoming), Virtual Private GW in VPC to Customer GW in on-prem, **Route Propagation** option in vpc needs to be enabled, to ping EC2 in vpc enable icmp in security group

- **Direct Connect** is regional, dedicated private connection from Virtual Private GW in VPC to on-prem, no security by default, add ipsec vpn for security
- Direct Connect is expensive and one month lead time, dedicated connection (1/10/100 Gbps), hosted connection (50/500 Mbps, 10 Gbps), supports ipv4/ipv6

- **Transit GW** is regional, can work cross region, connect multiple VPCs to Site-to-site VPN or Direct Connect to on-prem
- Transit Gateway can help in quickly adding VPCs, accounts, VPN capacity, or Direct Connect gateways to meet unexpected demand
- Transit GW **ONLY** option for multicast, can have blackhole route to drop traffic, creates two separate tunnels (outgoing, incoming) 1.25 Gbps each, cost by GB of data

- VPC Flow Logs allows capture of ip traffic in/out, can be enabled at vpc/subnet/eni level, can be sent to s3/cloudwatch logs, for debugging NACL/Security Group issues or connectivity issues
- Traffic Mirroring allows one to capture and inspect network traffic in a VPC in a non-intrusive way

IAM

- IAM has global scope and controls who is authenticated (identity) and who is authorized (has access) to use resources
- IAM first user is the **ROOT** and has full control over the account and its resources and **CANNOT** be restricted
- IAM users can be organized into **Groups** AND groups **CANNOT** contain other groups
- IAM **Role** is an identity with specific permissions that can used for delegation OR assumed by users/services
- IAM **Policy** is a JSON document that allows one to define permissions that explicitly allow/deny access to resources
- IAM explicit **DENY** overrides everything, explicit **ALLOW** in identity/resource policy overrides
- IAM **Permissions Boundary** is a type of policy that sets the maximum permissions that an identity (users OR roles NOT groups) can be granted
- IAM **Credentials Report** is an account-level report that lists all the user accounts and the status of their credentials
- IAM **Access Advisor** is an user-level report that shows the service permissions granted to a user and when those services were last accessed

- **Organizations** is an account management service that enables one to consolidate multiple accounts for better account/billing management to meet budgetary, security, and compliance needs
- Organizations main account is called the **Management Account** and has full admin power

- Organizations enables one to automate the creation of member accounts programmatically, get volume discounts for some services that have tiered pricing
- **Service Control Policy (SCP)** applies to organization/accounts level, **NOT** for management account, AND used to restrict access privileges (must have explicit allow)

- SCP only applies to users/roles in the organization or account (**NOT** service level roles)

Route 53

- Route 53 is 100% available DNS service, can handle domain registration/traffic routing/health checking
- When a domain is registered, it creates a Zone File (Hosted Zone) that costs 0.50 per month
- Zone File contains all the DNS records which map the domain names to target values
- DNS hostnames and DNS resolution are required settings for private hosted zones
- DNS A record maps a domain name to an IPv4 address AND AAAA record maps a domain name to an IPv6 address
- CNAME record must have an associated A or AAAA record, source cannot be root domain, can point to domain hosted anywhere, can be used to map one domain name to another, incurs charges for queries
- Alias record is custom extension to DNS, source can be root domain, source cannot be EC2, TTL auto set and cannot be changed, can route traffic to selected AWS resources (CloudFront/S3/ALB/API Gateway/Global Accelerator), can route traffic from one record in a hosted zone to another record, no charge for queries

- Health checks for load balancer or other public resources, health checkers around the globe, can monitor for cloudwatch alarms

- Simple Routing Policy - for single target resource, for multiple targets in A record then a random one is picked, **NO** health checks
- Weighted Routing Policy - for multiple target resources with percentage to each specific target, all targets same DNS record name/type, a percentage of zero means stop sending that target, useful for LB between regions and testing new app version, supports health checks
- Latency Routing Policy - for targets in multiple regions, latency based on traffic between client and region, supports health checks
- Failover Routing Policy - for implementing active-passive failover strategy across regions, only primary and secondary, health check a **MUST**
- Geolocation Routing Policy - for routing based on user location (continent/country/state), supports health checks
- Geoproximity Routing Policy - for routing traffic based on location of resources and optionally shift traffic based on a bias value (more value between 1 to 99 means more traffic, lesser value between -1 to -99 means lesser traffic), traffic flow **MUST** be enabled
- IP-Based Routing Policy - for routing based on list of ip cidrs of users
- Multivalue Routing Policy - for routing to one of the healthy targets of the many picked at random

S3

- S3 is global service, with 11 9s durability across multiple AZs

- S3 max object size 5 TB, bucket at region scope, bucket name **MUST** be globallly unique across all regions
- S3 buckets flat structure (no hierarchy), mimic folder using prefix, strong read-after-write consistency, multipart a **MUST** for > 5 GB,
- S3 **Transfer Acceleration** enabled at bucket level, uses CloudFront to speed transfer in the region, **ONLY** pay for accelerated data transfer
- S3 byte-range fetches for parallel fetch, better for retries (on errors)
- S3 object is owned by the account that uploaded it by default, even if the bucket is owned by another account
- S3 bucket by default private, identity-based for users/groups/roles/resources at bucket/object level
- S3 resource-based bucket policy for bucket level (applies to all objects) AND for cross account access
- S3 **Access Points** are named network endpoints attached to buckets, simplify access management for services to perform object operations
- S3 access points work in conjunction with bucket policy, can be used to grant permissions based on bucket prefix, can be configured to only allow from VPC
- S3 bucket can host static website accessible from the Internet, only include static web content as individual webpages with client-side scripts
- S3 static website bucket must enable **CORS** for clients to make cross-origin requests
- S3 versioning at bucket level, once enabled **CANNOT** be disabled **ONLY** suspended, delete of object inserts delete marker
- S3 deletion of specific version permanent, MFA delete enabled at bucket **ONLY** by owner using CLI, MFA delete prevents accidental deletes
- S3 replication at bucket, versioning a **MUST** for source AND target, IAM role for cross region, async for cross region
- S3 delete marker replication **MUST** enable a flag, delete of specific version OR delete markers **NOT** replicated to target

- S3 Standard for frequent access, min storage 30 days before transition to other tiers, can sustain loss of 2 facilities, useful for big data, mobile/gaming, content distribution
- S3 Standard-Infrequent Access for less-frequent/rapid access, min storage 30 days, charged for retrieval
- S3 Intelligent Tiering automatically moves between access tiers based on usage
- S3 One Zone-Infrequent Access for less frequently accessed reproducible/derived data, stores objects in only one AZ, min storage 30 days
- S3 Glacier Instant Retrieval allows millisecs retrieval, min storage 90 days, charged for retrieval, for data accessed per quarter
- S3 Glacier Flexible Retrieval allows mins to hours retrieval, min storage 90 days, charged for retrieval, for data archival
- S3 Glacier Deep Archive retrieval in hours (upto 48 hrs), min storage 180 days, for compliance/regulatory needs

- S3 Lifecycle Rules with versioning at bucket level, rules on prefix and object tags, no transition from other tiers to Standard
- S3 Lifecycle Rules no transition from Intelligent Tiering/One Zone-IA to Standard-IA, exipration can delete old versions after a period of time

- S3 Select/Glacier Select allows SQL expressions to select from large archive (zip) in a bucket, reduces network transfer cost

- S3 Server-Side Encryption with S3 Managed Keys (SSE-S3) enabled by default for new buckets/objects
- S3 Server-Side Encryption with Key Management Service (KMS) Managed Keys (SSE-KMS) has a default KMS key associated with S3, key usage audit via CloudTrail
- S3 Server-Side Encryption with Customer-Provided Keys (SSE-C) customer responsible for managing/providing encryption key, encryption key **NOT** stored by S3

- S3 bucket that is configured to host a static website **MUST** have the same name as the domain or subdomain

- S3 Pre-Signed URL allows temporary access to S3 objects for external customers, have an expiry duration

- S3 Access Logs enabled for audit purposes, **MUST** create a separate Logging bucket (with S3 Log Delivery permission)

- S3 Event Notifications event for create/delete/restore/replicate can be sent to sns/sqs/lambda, can integrate with event bridge for wider use-case, can be used to create thumbnails of images

- S3 Object Lambda for get requests via lambda, need to enable S3 Lambda Access Point in addition to bucket policy, useful for removing sensitive info (PII) from objects OR for decompression of object (bzip2, gzip, snappy, zlib, etc)

- S3 Object Lock can help prevent objects from being deleted or overwritten for a fixed amount of time or indefinitely
- S3 Object Lock with governance retention mode means some users can still overwrite/delete with IAM perms
- S3 Object Lock with compliance retention mode **NO** one can overwrite/delete
- S3 bucket with Object Lock enabled, one **CANNOT** disable object lock or suspend versioning for that bucket

CloudFront and Global Accelerator

- CloudFront is a global CDN service that speeds up distribution of static OR dynamic web content to users via Edge Locations with the lowest latency
- CloudFront uses a 24 hr default TTL with support for cache invalidation, can associate different TTLs for different content types
- CloudFront integration with S3 (origin) for static data distribution, enhanced access controls (origin access controls in addition to bucket policy), data uploads
- Cloudfront can also have custom http backend (ALB to private EC2, public EC2, S3 static website, any http location)
- CloudFront data not cached is pulled from the origin via the AWS global network (**NOT** internet)
- Cloudfront can be configured to go to specific origin based on path pattern

- CloudFront signed URL provides limited access to content for customers without cookie support, can specify validity date/time AND ip addr
- CloudFront signed cookies allows one to control access to multiple member-only private files (set-cookie for members only)
- CloudFront has support for imposing geo (country) restrictions using geo IP database via allowlist/blocklist
- Cloudfront data out cost (in GB) varies based on region

- Global Accelerator is a networking service for predictable app performance by routing traffic through the AWS global network (**NOT** internet)
- Global Accelerator creates two static anycast IP addresses that is used for intelligent routing
- Global Accelerator works with Elastic IP/EC2/ALB/NLB, can either be public/private, provides automatic DDoS protection via Shield

ECS and EKS

- ECS is a managed proprietary container orchestration service for deploying/managing/scaling containerized apps (ECS tasks)
- ECS support linux/windows, allows one to attach an ALB/NLB in front of cluster of ECS tasks (running container instances)
- ECS supports two types of launch environments - **EC2 launch type** AND **Fargate launch type** (serverless)
- EC2 launch type includes an **ECS Container Agent** that exposes APIs to gather info about the container instance
- ECS Optimized AMI includes the ECS Container Agent by default
- Fargate launch type is serverless and charges are based on per running task
- ECS service defines how to run tasks with auto scaling and desired count, while ECS cluster is a logical grouping of services that can run across AZs in region
- ECS auto scaling uses CloudWatch Metrics to trigger scaling based on demand
- ECS auto scaling has two categories - ECS service auto scaling (for EC2 launch type and Fargate launch type) AND ECS cluster auto scaling (only for EC2 launch type)
- ECS service auto scaling types - Target Tracking, Step Scaling, Scheduled Scaling
- ECS Data Volumes - can be EFS for multi AZ AND for EC2/Fargate launch types, EBS only for EC2 launch type
- ECS networking mode AWSVPC allocates separate ENI (ip addr, security group) to each task which helps with separate security policies for each task

- ECR is a managed proprietary container image registry with support for OCI/Docker registry API AND uses S3 as the image store
- ECR supports private/public repositories, private repository for customer images only (not public and controlled via IAM access controls)
- ECR Image Scanning helps identify vulnerabilities in customer images

- EKS is a managed open source Kubernetes service (which can also be run on-prem) with two types of launch environments - **EC2 launch type** AND **Fargate launch type** (serverless)
- EKS has two types of EC2 launch types - managed node groups (AWS managed) and self-managed nodes with BOTH using the auto scale group for scaling

- EKS Cluster consists of an EKS Control Plane and a group of EKS worker nodes (grouping of EKS Pods)
- EKS Cluster Auto Scaling of two types - Vertical Pod Autoscaler (adjusts the CPU and Memory reservations for the EKS pods), Horizontal Pod Autoscaler (scales the number of EKS pods in the Replica Set based on the resource CPU utilization)
- EKS supports both ALB and NLB for load balancer
- EKS Data Volumes - can be EFS for EC2/Fargate launch types, EBS only for EC2 launch type

Lambda

- Lambda is a managed serverless compute service that can be used for event-driven architecture to execute code in response to receiving events
- Lambda supports modern programming languages (Node.js/Python/Java 8 and above/C# .Net Core/PowerShell/Golang/Ruby)
- Lambda function **MUST** have short code execution for up to 15 mins
- Lambda **SYNC** execution triggered by CLI, SDK, API
- Lambda **ASYNC** execution can be triggered by S3/SQS/SNS/CloudWatch/DynamoDB Streams/Kinesis Data Streams AND must be idempotent with 3 retries on error
- Lambda functions execute concurrently (in parallel) based on the incoming events AND if reaches a concurrency limit (based on region), then throttles the execution
- Lambda charged based on compute usage (in millisecs), one can provision more RAM (up to 10 GB)
- Lambda function can be configured to run in a custom VPC by providing the VPC ID, the subnets to use, and attach the Security Groups
- Lambda function can be configured to be invoked from within RDS database (RDS PostgreSQL or Aurora MySQL)
- Lambda SnapStart takes a snapshot of memory/disk of an initialized java function (11 an above) and caches it for faster future processing
- Lambda @Edge code can be writtein in Node.js or Python AND can be attached to CloudFront (in us-east-1 for global) for complex request/response transformation (CloudFront Function can only manipulate header/cookies/url and validate JWT)
