Questions and AWS resource differences

AWS Lambda

Can have a maximum runtime of only 15 minutes.

A great option for hosting microservices that does a certain task independently of others.

Combine with API Gateway to host the API, they both scale very fast and is great against DDos. Hard to overwhelm these services.

AWS API Gateway

Built in protection against DDOS.

AWS System Manager Session Manager

Using it to connect to your private instance say EC2 comes at no extra charge.

AWS S3

If the items that you are uploading reaches above 100MB, then you should generally consider using multipart upload, to split each files up into smaller parts, upload the smaller parts and then merge them together after they are uploaded into a complete object.

Multipart helps with easily resume failed upload attempts.

`aws s3 cp` command automatically performs multipart uploading for you already.

AWS Peering

It has a limitation of 125 peer connections for each VPC.

AWS PrivateLink

Best used when you have client-server setup where you want to allow one or more consumer VPCs unidirectional access to a specific service (or a set of instances) in the service-provider VPC.

Here are the instructions on setting one up:

1. Create a network load balancer
2. Create service-consumer role in IAM
3. Set up an endpoint connection in the shared services VPC, and set it to automatically accept
4. Create consumer endpoints in each VPC that must access the shared VPC
5. Then point the network load balancer in the shared service VPC

Amazon ElastiCache

A cache-as-a-service. It let you deploy, manage, and scale a distributed in-memory cache in the cloud.

It is built to boost performance of web-based application by reducing database load through quick retrieval of data from cache.

Amazon DynamoDB

DynamoDB is a key-value database, scales automatically, and handle spiky access pattern.

DynamoDB also have DynamoDB streams, it can stream data that's in the database.

Amazon Kinesis

You can configure kinesis to delivery streaming data from a database and you can deposit the data into a S3 data lake

DocumentDB

It is not a key-value database, it stores the data in a JSON format

AWS IAM Role

Identity and access management is how you give a fine-grained access control on AWS resources, and control what kind of permission each of the resources have.

IAM role contains many IAM policies which is how the IAM role actually gain permissions, controlling what they can or cannot do in AWS. Policies are attached to each IAM role.

When you are giving permission to a IAM role with policies, always use least-privilege permissions which means you only give the permission that the AWS resources need to perform the task, you never give more than you need to.

Best practice for root user: Use multi-factor authentication + use a complex password for root user login. Don't make access key (OAuth)! It gives access to all of the AWS resources which is bad if someone get their hands on this access key.

File storage vs Object storage

File storage: Organized into tree-like hierarchy with directories and sub-directories. Just like a filesystem in Linux system, the root directory and so on.

For file storage, in order to access a particular file, you have to provide a path for it.

Object storage: No directories, all the files are stored into basically one "directory", and each file will have a unique identifier to make it easier to find that particular file. It can store lots amount of objects due to no limitation and is scalable.

File storage will have a faster latency since if you have the path to the particular file, it can locate it quickly. Object storage on the other hand is created with cost efficiency and scalability in mind, so those benefits comes with the cost of speed and performance.

Revision #6
Created 2023-01-04 22:29:34 UTC by Tamarine
Updated 2023-02-06 18:24:02 UTC by Tamarine