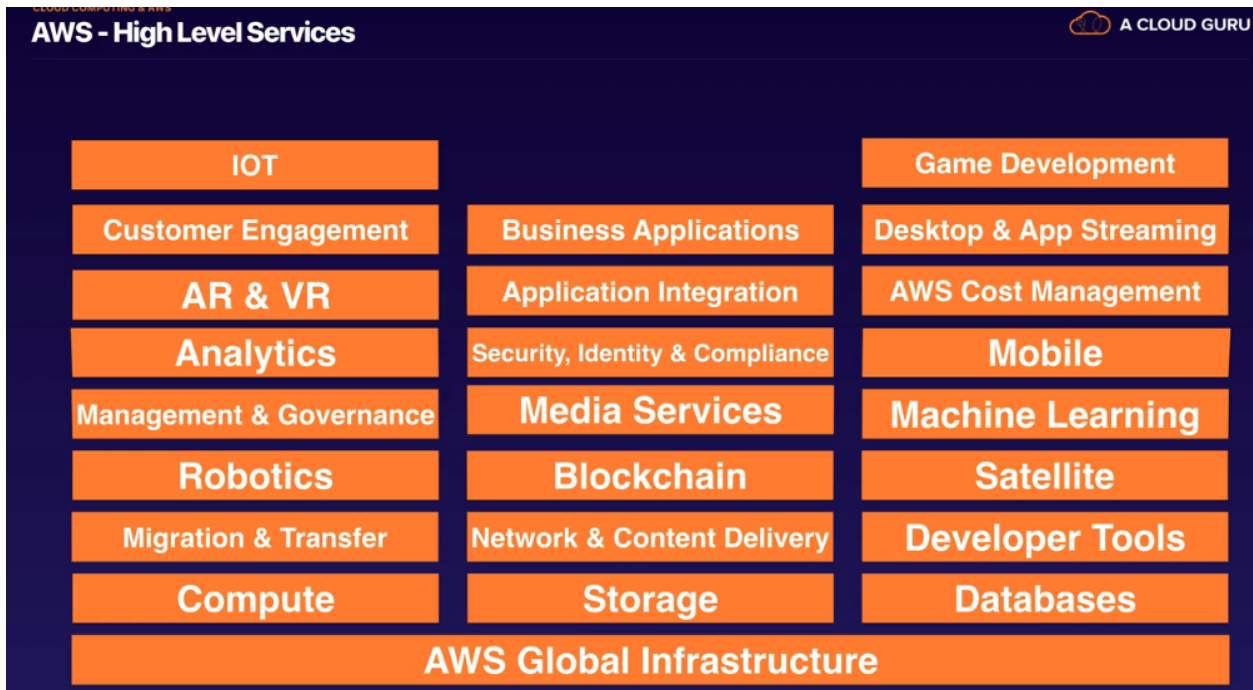


- Section 2
 - S3 Buckets
 - IAM
 - Load Balancing
 - Availability Zones
 - Domain Registration
 - AWS Billing
 - Security and Compliance
1. Exam blue print
 - 1.1. Cloud Concept 26%
 - 1.1.1. CC and Tech in one section
 - 1.2. Security 25 %
 - 1.2.1. Understand cloud concept and tech and some billing
 - 1.3. Technology 33%
 - 1.4. Billing and Pricing 16%
 2. Cloud computing and AWS
 - 2.1. 6 advantages of Cloud computing
 - 2.1.1. Trade Capital Expense for Variable Expense
 - 2.1.2. Benefit form Massive economies of scale
 - 2.1.3. Stop Guessing about capacity
 - 2.1.4. Increase speed and agility
 - 2.1.5. Stop spending money running and maintaining data centers
 - 2.1.6. Go global in minutes
 - 2.2. 3 different type of cloud computing
 - 2.2.1. Infrastructure As A Service (IAAS)
 - 2.2.1.1. Servers (physical or virtual), Operating system
 - 2.2.1.1.1. EC2
 - 2.2.2. Platform As A Service (PAAS)
 - 2.2.2.1. Manages Hardware and Operating systems, security, patching, updates and maintenance etc
 - 2.2.2.1.1. Elastic beanstalk (provision code) and AWS will take care of the hardware specs
 - 2.2.2.1.2. Amazon light sale
 - 2.2.3. Software As A Service (SAAS)
 - 2.2.3.1. Example
 - 2.2.3.1.1. gmail
 - 2.3. 3 different deployment types of cloud computing.
 - 2.3.1. Public cloud
 - 2.3.1.1. AWS
 - 2.3.1.2. Azure
 - 2.3.1.3. GCP
 - 2.3.2. Hybird
 - 2.3.2.1. Mix of public and private

2.3.3. Private Cloud (Or On Premise)

2.3.3.1. You manage it, your datacenter, openstack, Vmware, Hyper-V



Services relevant to AWS Cloud practitioner



What do you need to know to pass the Certified Cloud Practitioner Exam?

AWS Cost Management

Security, Identity & Compliance

Compute

Storage

Databases

AWS Global Infrastructure

- Certified Cloud Practitioner Exam
 - AWS Global Infrastructure
 - Compute
 - EC2
 - VM's in the cloud
 - Lambda
 - Code focused level up from EC2
 - Storage
 - Simple Storage Service (S3)
 - Glacier
 - Databases
 - Relational Database Service (RDS)
 - DynamoDB (Non Relational Databases)
 - Security, Identity & Compliance
 - Network
 - VPC
 - Route53
 - Amazon's net DNS service
 - Around the world AWS
 - Regions (about 25 regions as of 2019)
 - 2 or more zones (72 as of 2019)
 -
 - Gov cloud
 - Edge location (endpoints for AWS for caching content)
 - Cloud Front, Amazons content Delivery Network (CDN)
 - 150 location as of 2019
 - Exam Tips

- A Region is a physical location in the world which consists of two or more Availability Zones (AZ's)
- An AZ is one or more discrete data centers, each with redundant power, networking and connectivity, housed in separate facilities
- Edge Locations are endpoints for AWS which are used for caching content typically this consists of CloudFront, Amazon's Content Delivery Network (CDN)
- Choosing AWS Region (Particular region, and particular service)
 - Data Sovereignty Laws
 - Latency to end users
 - AWS services (US-EAST1 primary region)
- Exam Tips
 - Basic- Free
 - Developer- \$29 a month (scales based on usage)
 - Business- \$100 a month (scales based on usage)
 - Enterprise - \$15,000 a month (scales based on usage) - TAM, Technical Account Manager)

<https://shahraizq.signin.aws.amazon.com/console>

Shahraiz_e

Access key ID: AKIAX4UUXULER7JXWBBE

Secret access key: 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

Pass: +1Gv^&UlbJ48-4(

- Exam Tips
 - IAM Identity Access Management. It is Global, you do not specify a region when dealing with IAM. When you create a user or group, this is created GLOBALLY.
 - AWS Platform in 3 ways
 - Via the Console
 - Programmatically (Using the Command Line)
 - Software Developers Kit (SDK)
 - Group is a place to store your users
 - Users inherit all the permissions the group they're in has
 - Group examples
 - Developers
 - System admin
 - Human resources,
 - Finance etc
 - Permissions in the group
 - To set permissions in a group, apply a policy to that group
 - Policies consist of Java Script Object Notation (JSON)
 - Policies have Key value pair

- Key and value

- S3 (Simple Storage service) (safe place to store your files)
 - Provides developers and IT teams w/ secure, durable, highly-scalable object storage
 - Easy to use w/simple web services interface to store and retrieve any amount of data from anywhere on the web.
 - Place for flat files(files that don't change)
 - Object based storage (oposite to block based storage (install operating system, or database etc)
 - The data is spread across multiple devices and facilities
 - S3 Basics
 - Object- based
 - Files can be 0 bytes to 5 TB
 - Unlimited storage
 - Files stored in buckets
 - Universal namespace- name must be unique
 - Url always starts w/ s3(service), region which the bucket is located (eu-west-1(Ireland)), amazonaws.com, /bucketname
<https://s3-eu-west-1.amazonaws.com/acloudguru>
 - Creates a DNS entry so the bucket name must be unique globally.
 - When file has been uploaded to S3, you will receive HTTP 200 code if upload was successful.
 - S3 is key value store
 - S3 is object based objects example below
 - Key (name of object)
 - Value(the data, made up of sequence of bytes
 - Version ID (important for versioning)
 - Metadata(data about data stored)
 - Subresources
 - Access control lists
 - Torrent
- Data Consistency Model for S3
 - Read after write consistency for PUTS of new Objects
 - If you write a new file and read it immediately afterwards, you will be able to view that data.
 - Eventual Consistency for overwrite PUTS and DELETES (can take some time to propagate)
 - If you update AN EXISTING file or delete a file and read it immediately, you may get the older version, or you may not. Basically changes to objects can take a little bit of time to propagate.

- S3 - Guarantees
 - Built for 99.99% availability
 - Amazon Guarantee 99.9 availability
 - Amazon guarantees 99.999999999% durability for S3 information (Remember 11 x 9s)
 - 11 9s of durability

- S3 - Features
 - Tiered Storage Available
 - Lifecycle Management
 - Versioning
 - Encryption
 - Secure you data using Access Control Lists and Bucket Policies
 - Access Control lists
 - File or object level
 - Bucket Policies
 - Bucket level

- Amazon S3 offers a range of storage classes designed for different use cases. These include **S3 Standard** for general-purpose storage of frequently accessed data; **S3 Intelligent-Tiering** for data with unknown or changing access patterns; **S3 Standard-Infrequent Access (S3 Standard-IA)** and **S3 One Zone-Infrequent Access (S3 One Zone-IA)** for long-lived, but less frequently accessed data; and **Amazon S3 Glacier (S3 Glacier)** and **Amazon S3 Glacier Deep Archive (S3 Glacier Deep Archive)** for long-term archive and digital preservation. If you have data residency requirements that can't be met by an existing AWS Region, you can use the **S3 Outposts storage** class to store your S3 data on-premises. Amazon S3 also offers capabilities to manage your data throughout its lifecycle. Once an S3 Lifecycle policy is set, your data will automatically transfer to a different storage class without any changes to your application.
 - S3 Storage classes
 - S3 Standard
 - S3 Standard offers high durability, availability, and performance object storage for frequently accessed data. Because it delivers low latency and high throughput, S3 Standard is appropriate for a wide variety of use cases, including cloud applications, dynamic websites, content distribution, mobile and gaming applications, and big data analytics. S3 Storage Classes can be configured at the object level and a single bucket can contain objects stored across S3 Standard, S3 Intelligent-Tiering, S3 Standard-IA, and S3 One Zone-IA. You can also use S3 Lifecycle policies to automatically transition objects between storage classes without any application changes.
 - Data is stored in multiple locations. So it is resilient against events that affect an entire **Availability Zone**
 - Most expensive storage class among all others.
 - Low latency and high throughput performance (High Availability)

- Designed for durability of 99.999999999% of objects across multiple Availability Zones
 - Resilient against events that impact an entire Availability Zone
 - Designed for 99.99% availability over a given year
 - Backed with the [Amazon S3 Service Level Agreement](#) for availability
 - Supports SSL for data in transit and encryption of data at rest
 - S3 Lifecycle management for automatic migration of objects to other S3 Storage Classes
- Unknown or changing access
 - S3 - Intelligent Tiering
 - [Amazon S3 Intelligent-Tiering \(S3 Intelligent-Tiering\)](#) is the only cloud storage class that delivers automatic storage cost savings when data access patterns change, without performance impact or operational overhead. For a small monthly object monitoring and automation charge, S3 Intelligent-Tiering monitors access patterns and automatically moves objects that have not been accessed to lower cost access tiers. S3 Intelligent-Tiering delivers automatic storage cost savings in two low latency and high throughput access tiers. For data that can be accessed asynchronously, customers can choose to activate automatic archiving capabilities within the S3 Intelligent-Tiering storage class. There are no retrieval charges in S3 Intelligent-Tiering. If an object in the infrequent access tier is accessed later, it is automatically moved back to the frequent access tier. No additional tiering charges apply when objects are moved between access tiers within the S3 Intelligent-Tiering storage class.
 - S3 Intelligent-Tiering is the ideal storage class for data with unknown, changing, or unpredictable access patterns, independent of object size or retention period. You can use S3 Intelligent-Tiering as the default storage class for data lakes, analytics, and new applications.
 - S3 Intelligent-Tiering works by storing objects in access tiers: two low latency access tiers optimized for frequent and infrequent access, and two opt-in archive access tiers designed for asynchronous access that are optimized for rare access. Objects uploaded or transitioned to S3 Intelligent-Tiering are automatically stored in the Frequent Access tier. S3 Intelligent-Tiering monitors access patterns and then moves objects that have not been accessed in 30 consecutive days to the Infrequent Access tier. Once you have activated one or both of the Archive Access tiers, S3 Intelligent-Tiering will move objects that haven't been accessed for 90 consecutive days to the Archive Access tier and then after 180 consecutive days of no access to the Deep Archive Access tier. If the objects are accessed later, S3 Intelligent-Tiering moves the objects back to the Frequent Access tier. If the object you are retrieving is stored in the Archive or Deep Archive tiers, before you can retrieve

the object you must first restore a copy using `RestoreObject`. For information about restoring archived objects, see [Restoring Archived Objects](#).

- - Low latency and high throughput performance
 - Automatically moves the data between two access tiers. (Infrequent Access and Frequent Access)
 - Small monthly monitoring and auto-tiering fee
 - Automatically optimizes storage costs for data with changing access patterns
 - Stores objects in two access tiers, optimized for frequent and infrequent access
 - Frequent and Infrequent Access tiers have same low latency and high throughput performance of S3 Standard
 - Activate optional automatic asynchronous archive capabilities for objects that become rarely accessed
 - Archive access and deep Archive access tiers have same performance as Glacier and Glacier Deep Archive
 - Designed for durability of 99.999999999% of objects across multiple Availability Zones
 - Designed for 99.9% availability over a given year
 - Small monthly monitoring and auto-tiering charge
 - No operational overhead, no retrieval charges, no additional tiering charges apply when objects are moved between access tiers within the S3 Intelligent-Tiering storage class
 - No minimum storage duration
- Infrequent Access
 - S3- Standard IA (Infrequent Access)
 - S3 Standard-IA is for data that is accessed less frequently, but requires rapid access when needed. S3 Standard-IA offers the high durability, high throughput, and low latency of S3 Standard, with a low per GB storage price and per GB retrieval charge. This combination of low cost and high performance make S3 Standard-IA ideal for long-term storage, backups, and as a data store for disaster recovery files. S3 Storage Classes can be configured at the object level and a single bucket can contain objects stored across S3 Standard, S3 Intelligent-Tiering, S3 Standard-IA, and S3 One Zone-IA. You can also use S3 Lifecycle policies to automatically transition objects between storage classes without any application changes.
 - High Availability and Low Latency (Same as S3 Standard)
 - Offers greater availability and resiliency than the OneZone-IA storage.

- The durability of 99.999999999% and availability of 99.99% availability over a given year
 - Less expensive than S3 Standard storage but you will be charged a retrieval fee hence suitable for infrequently accessed data.
 - Same low latency and high throughput performance of S3 Standard
 - Designed for durability of 99.999999999% of objects across multiple Availability Zones
 - Resilient against events that impact an entire Availability Zone
 - Data is resilient in the event of one entire Availability Zone destruction
 - Designed for 99.9% availability over a given year
 - Supports SSL for data in transit and encryption of data at rest
 - S3 Lifecycle management for automatic migration of objects to other S3 Storage Classes
- S3 One Zone - IA
 - S3 One Zone-IA is for data that is accessed less frequently, but requires rapid access when needed. Unlike other S3 Storage Classes which store data in a minimum of three Availability Zones (AZs), S3 One Zone-IA stores data in a single AZ and costs 20% less than S3 Standard-IA. S3 One Zone-IA is ideal for customers who want a lower-cost option for infrequently accessed data but do not require the availability and resilience of S3 Standard or S3 Standard-IA. It's a good choice for storing secondary backup copies of on-premises data or easily re-creatable data. You can also use it as cost-effective storage for data that is replicated from another AWS Region using S3 Cross-Region Replication.
 - S3 One Zone-IA offers the same high durability†, high throughput, and low latency of S3 Standard, with a low per GB storage price and per GB retrieval charge. S3 Storage Classes can be configured at the object level, and a single bucket can contain objects stored across S3 Standard, S3 Intelligent-Tiering, S3 Standard-IA, and S3 One Zone-IA. You can also use S3 Lifecycle policies to automatically transition objects between storage classes without any application changes.
 - - Low Latency and High throughput performance
 - The durability of 99.999999999% and availability of **99.5% availability** over a given year
 - Data will be lost if the Availability Zone where the data is stored is destroyed.
 - Suitable for larger objects greater than 128 KB kept for at least 30 days (charged minimum for 30 days)

- Same low latency and high throughput performance of S3 Standard
 - Designed for durability of 99.999999999% of objects in a single Availability Zone†
 - Designed for 99.5% availability over a given year
 - Supports SSL for data in transit and encryption of data at rest
 - S3 Lifecycle management for automatic migration of objects to other S3 Storage Classes
 - † Because S3 One Zone-IA stores data in a single AWS Availability Zone, data stored in this storage class will be lost in the event of Availability Zone destruction.
- Archive
 - S3 Glacier
 - S3 Glacier is a secure, durable, and low-cost storage class for data archiving. You can reliably store any amount of data at costs that are competitive with or cheaper than on-premises solutions. To keep costs low yet suitable for varying needs, S3 Glacier provides three retrieval options that range from a few minutes to hours. You can upload objects directly to S3 Glacier, or use S3 Lifecycle policies to transfer data between any of the S3 Storage Classes for active data (S3 Standard, S3 Intelligent-Tiering, S3 Standard-IA, and S3 One Zone-IA) and S3 Glacier. For more information, visit the [Amazon S3 Glacier page](#) » S3 Glacier is a low-cost storage class for data archiving where data access is infrequent. It provides a configurable retrieval time for the data from minutes to hours. This storage class uses a very low-cost Glacier storage service but the objects are still managed through S3.
 - Low-cost design for long-term archiving
 - Data will be available in case of entire Availability Zone destruction
 - The durability of 99.999999999% and availability of 99.9% availability over a given year
 - It has a minimum storage duration period of 90 days.
 - Designed for durability of 99.999999999% of objects across multiple Availability Zones
 - Data is resilient in the event of one entire Availability Zone destruction
 - Supports SSL for data in transit and encryption of data at rest
 - Low-cost design is ideal for long-term archive
 - Configurable retrieval times, from minutes to hours
 - S3 PUT API for direct uploads to S3 Glacier, and S3 Lifecycle management for automatic migration of objects
 - S3 Glacier Deep Archive
 - S3 Glacier Deep Archive is Amazon S3's lowest-cost storage class and supports long-term retention and digital preservation for data that may be accessed once or twice in a year. It is designed for customers

— particularly those in highly-regulated industries, such as the Financial Services, Healthcare, and Public Sectors — that retain data sets for 7-10 years or longer to meet regulatory compliance requirements. S3 Glacier Deep Archive can also be used for backup and disaster recovery use cases, and is a cost-effective and easy-to-manage alternative to magnetic tape systems, whether they are on-premises libraries or off-premises services. S3 Glacier Deep Archive complements Amazon S3 Glacier, which is ideal for archives where data is regularly retrieved and some of the data may be needed in minutes. All objects stored in S3 Glacier Deep Archive are replicated and stored across at least three geographically-dispersed Availability Zones, protected by 99.999999999% of durability, and can be restored within 12 hours. The S3 Glacier Deep Archive provides the **lowest-cost storage** class and supports long-term retention and digital preservation for data that may be accessed only once or twice in a year. It is ideal for those industries which store data for 5-10 years or longer like healthcare, finance, etc. It can also be used for backup and disaster recovery.

- Lowest cost storage option in S3
- The durability of 99.999999999% and availability of 99.9% availability over a given year
- Retrieval costs can be reduced by using bulk retrieval
- It has a minimum storage duration period of 180 days
- Designed for durability of 99.999999999% of objects across multiple Availability Zones
- Lowest cost storage class designed for long-term retention of data that will be retained for 7-10 years
- Ideal alternative to magnetic tape libraries
- Retrieval time within 12 hours
- S3 PUT API for direct uploads to S3 Glacier Deep Archive, and S3 Lifecycle management for automatic migration of objects
-
- S3 on Outposts
 - S3 Outposts storage class
 - Amazon S3 on Outposts delivers object storage to your on-premises AWS Outposts environment. Using the S3 APIs and features available in AWS Regions today, S3 on Outposts makes it easy to store and retrieve data on your Outpost, as well as secure the data, control access, tag, and report on it. S3 on Outposts provides a single Amazon S3 storage class, named S3 Outposts, which uses the S3 APIs, and is designed to durably and redundantly store data across multiple devices and servers on your Outposts. S3 Outposts storage class is ideal for workloads with local data residency requirements, and to satisfy demanding performance needs

- S3 Charges
 - Charged for S3 in the following ways
 - Storage
 - Requests
 - Storage Management Pricing
 - Data Transfer Pricing
 - Transfer Acceleration
 - Cross Region Replication Pricing
- S3 Transfer Acceleration
 - Enables fast, easy, and secure transfers of files over long distances between you end users and an S3 bucket.
 - Takes advantage of Amazon CloudFronts globally distributed edge locations. As the data arrives at an edge location, data is routed to Amazon S3 over an optimized network path.
- Cross Region Replication
 - Copies everything to another region, use case is disaster recovery
- Exam tips
 - S3 is object-based (allows you to upload files) (not suitable to install an OS)
 - Files can be from 0 Bytes to 5 TB
 - There is unlimited storage
 - Files are stored in Buckets
 - S3 is a universal namespace, means it must be unique
 - <https://s3-eu-west-1.amazonaws.com/acloudguru>
 - Successful uploads will generate a HTTP 200 status code
- Key Fundamentals of S3
 - Key(name of the object) Value (the data and made up of sequence of bytes) pair
 - S3 model
 - Read after Write consistency for PUTS of new Objects
 - Eventual Consistency for overwrite PUTS and DELETES (Can take some time to propagate)
 - 7 different storage classes
 - S3 standard
 - S3-IA (Infrequently Accessed)
 - Lower fees then S3
 - S3 One Zone- IA
 - S3 - Intelligent Tiering
 - Uses Machine learning moves data around different classes
 - S3 Glacier
 - S3 Glacier Deep Archive
 - Data retrieval times of 12 hours
 - S3 Outposts
 - Deliver object storage on-premises AWS outpost environments
 - Route53
 - Amazon DNS service
- Restrict Bucket Access

Bucket Policies - applies across the whole bucket

- Object Policies - Applies to individual files
- IAM Policies to Users & Group - Applies to Users & Groups

- Serverless website on S3
- Use bucket policies to make entire S3 buckets public.
- Use S3 to host STATIC Websites (such as .html). Websites that require database connections such as wordpress etc cannot be hosted on S3.
- S3 scales automatically to meet your demand. Many enterprises will put static websites on S3 when they think there is going to be a large number of request (such as movie preview)

-Versioning w/S3

- 1. Store all version of an object
 - Including all writes and even if you delete an object
- 2. Great backup tool
- 3. Versioning cannot be disabled
 - Once enabled, versioning cannot be disabled - only suspended
- 4. Integrates with life cycles rules
- 5. Versioning's MFA Delete capability
 - Uses multi-factor authentication; can be used to provide an additional layer of security

-Cloud Front

- A content delivery network (CDN) is a system of distributed servers (network) that delivers web pages and other web content to a user based on the geographic locations of the user, the origin of the webpage, and a content delivery server.
 - Edge location: this is the location where content will be cached. This is separate to an aws Region/AZ.
 - Edge locations are not just READ only, you can write to them too.(ie put an object on to them)
 - Objects are cached for the life of the TTL (Time to live)
 - TTL: The Default TTL is now 86400 seconds (24 hrs). If the Minimum TTL is set to more than 86400 seconds, the Default TTL changes to the value of the Minimum TTL.

schedule that only requires a fraction of a day, a week, or a month.

- 3. Spot (bidding)
 - Enables you to bid whatever price you want for instance capacity, providing for even greater savings if your applications have flexible start and end times.
 - Applications THAT have flexible start and end times
 - Applications that are only feasible at very low compute prices
 - Users with urgent computing needs for large amounts of additional capacity
- 4. Dedicated Hosts
 - Physical EC2 servers dedicated for your use. Dedicated hosts can help you reduce costs by allowing you to use your existing server-bond software licences.
 - Useful for regulatory requirements that may not support multi-tenant virtualization.
 - Great for licensing which does not support multi-tenancy or cloud
 - Can be purchased on-demand (hourly)
 - Can be purchased as a Reservation for upto 70% off the On-Demand price.
- Convertible vs reservable instances.

Family	Speciality	Use case
F1	Field Programmable Gate Array	Genomics research, financial analytics, real-time video processing, big data etc
I3	High Speed Storage	NoSQL DBs, Data Warehousing etc
G4	Graphics Intensive	Video Encoding/ 3D Application Streaming
H1	High Disk Throughput	MapReduce-based workloads, distributed file systems such as HDFS and MapR-FS
T4g	Lowest Cost, General Purpose	Web Servers/Small DBs
D2	Dense Storage	Fileservers/Data Warehousing/Hadoop
R6g	Memory Optimized	Memory Intensive Apps/DBs
M6g	General Purpose	Application Servers
C6g	Compute Optimized	CPU Intensive Apps/DBs
P3	Graphics/General Purpose GPU	Machine Learning, Bit Coin Mining etc
X1e	Memory Optimized	SAP HANA/Apache Spark etc
Z1D	High compute capacity and a high memory footprint.	Ideal for electronic design automation (EDA) and certain relational database workloads with high per-core licensing costs.
A1	Arm-based workloads	Scale-out workloads such as web servers
U-12TB1	Bare Metal	Bare metal capabilities that eliminate virtualization overhead

EC2 Instance Types - Mnemonic

- **F** - For FPGA
- **I** - For IOPS
- **G** - Graphics
- **H** - High Disk Throughput
- **T** - Cheap general purpose (think T2 Micro)
- **D** - For Density
- **R** - For RAM
- **M** - Main choice for general purpose apps
- **C** - For Compute
- **P** - Graphics (think Pics)
- **X** - Extreme Memory
- **Z** - Extreme Memory AND CPU
- **A** - Arm-based workloads
- **U** - Bare Metal

- Amazon EBS : virtual disks in the cloud
 - Allows you to create storage volumes and attach them to Amazon EC2 instances. Once attached, you can create a file system on top of these volumes, run a database, or use them in any other way you would use a block device. Amazon EBS volumes are placed in specific Availability Zone, where they are automatically replicated to protect you from the failure of a single component.
 - SSD
 - General purpose SSD (GP2)- balances price and performance for a wide variety of workloads
 - Provisioned IOPS SSD (IO1)- Highest- performance SSD volume for mission-critical low-latency or high- throughput workloads
 - Magnetic
 - Throughput Optimized HDD (ST1) - Low cost HDD volume designed for frequently accessed, throughput-intensive workloads
 - Cold HDD (SC1) - Lowest cost HDD volume designed for less frequently accessed workloads (File Servers).
 - Magnetic - Previous Generation

- Exam tips: EC2- Virtual server in the Cloud
 - Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides resizable compute capacity in the cloud. Amazon EC2 reduces the time required to obtain and boot new server instances to minutes, allowing to quickly scale capacity, both up and down, as your computing requirements change.
 - 4 different pricing models
 - 1 On demand
 - Allows you to pay a fixed rate by the hour (or by the second) with no commitment
 - 2 Reserved
 - Provides you with a capacity reservation, and offer a significant discount on the hourly charge for an instance. Contract terms are 1 year or 3 year terms.
 - 3 Spot
 - Enables you to bid what ever price you want for instance capacity, providing for even greater savings if your applications have flexible start and end times.
 - 4 Dedicated Hosts
 - Physical EC2 servers dedicated for your use. Dedicated Hosts can help you reduce costs by allowing your to use your existing server-bond software licences.
- If the spot instance is terminated by Amazon EC2, you will not be charged for a partial hour of usage. However, if you terminated the instance yourself, you will be charged for any hour which the instance ran.

EC2 Instance Types - Mnemonic

- **F** - For FPGA
- **I** - For IOPS
- **G** - Graphics
- **H** - High Disk Throughput
- **T** - Cheap general purpose (think T2 Micro)
- **D** - For Density
- **R** - For RAM
- **M** - Main choice for general purpose apps
- **C** - For Compute
- **P** - Graphics (think Pics)
- **X** - Extreme Memory
- **Z** - Extreme Memory AND CPU
- **A** - Arm-based workloads
- **U** - Bare Metal

4 types

- SSD
 - General purpose SSD (GP2) - balances price and performance for a wide variety of workloads.
 - Provisioned IOPS SSD (IO1) - Highest - performance SSD volume for mission - critical low - latency or high- throughput workloads
- Magnetic
 - Throughput Optimized HDD (ST1) - Low cost HDD volume designed for frequently accessed, Throughput- intensive workloads.
 - Cold HDD (SC1) - Lowest cost HDD volume designed for less frequently accessed workloads (File Servers)
 - Magnetic - previous Generation

EC2 Lab

EC2: Amazon Elastic Compute Cloud (Amazon EC2) is just a virtual server (or servers) in the cloud. Amazon EC2 reduces the time required to obtain and boot new server instances to minutes, allowing you to quickly scale capacity, both up and down, as your computing requirements change.

Linux = SSH Port 22

Microsoft = RDP Port 3389

HTTP = Port 80

HTTPS = Port 443

CIDR Address range 0.0.0.0/0 Let every thing in
X.X.X.X/32 To let just one IP in

EC2 is compute based service. It is not serverless. Its a server!
Use a private key to connect to EC2

Security Groups are virtual firewalls in the cloud. You need to open ports in order to use them.
Popular ports are SSH(22), HTTP(80), HTTPS(443), RDP (3389)

Always design for failure. Have one EC2 instance in each availability zone.

Using the AWS command line -LAB

Exam tip

Users can interact with AWS in 3 ways

Using console

Command Line Interface (CLI)

Using the Software Development kits (SDKs)

Exam tips

Roles are much more secure then using access key id's and secret access keys and are easier to manage.

You can apply roles to EC2 instances at any time. When you do this the change takes place immediately.

Roles are universal. You do not need to specify what region they are in, similar to users.

Exam Tips

Load Balancers come in 3 different flavors.

Application Load Balancers, Network Load Balancers, Classic Load Balancers

Application Load Balancers - Layer 7 (Make Intelligent Decisions). Network Load Balancers - Extreme Performance/Static IP Addresses. Classic Load Balancers - Test & Dev, Keep Costs Low.

Databases101

Relational Databases

- Database
- Tables
- Rows
- fields

Relational Databases on AWS: RDS

- SQL Server
- Oracle
- MySQL Server
- PostgreSQL
- Aurora
- MariaDB

RDS VS Read Replicas

RDS has 2 key features:

- Multi-AZ-For Disaster Recovery
- Read Replicas - For Performance

Non Relational Database

- Collection
 - Table
- Document
 - Row
- Key Value Pairs
 - Fields
 - Nest for Flat
- The columns in the table can vary
- The will not affect other rows in the database
- A lot of flexibility

Amazon's Non-Relational Database is called DynamoDB

NoSQL or Non-relational DB and need performance = DynamoDB

OLTP vs OLAP

Online Transaction Processing (OLTP) differs from OLAP Online Analytics Processing (OLAP) in terms of the types of queries you will run

OLAP Online Transaction Processing	OLTP Online Analytics Processing
Net profit for EMEA and Pacific for the Digital Radio Product. Pulls in large numbers of records	Order number 2120121 pulls up a row of data such as Name, Date, Address to Deliver to, Delivery status etc.

Sum of Radios Sold in EMEA Sum of Radios Sold in Pacific Unit Cost of Radio in each region Unit Cost of Radio in each region Sales price of each radio Sales price - unit cost	
Data warehousing <ul style="list-style-type: none"> Do OLAP away from production DB Used for BI in tools like Cognos, JasperSoft, SQL Server Reporting Services, Oracle, Hyperion, SAP NetWeaver 	Takes/insert a row of data from the database
Used to pull VERY LARGE and complex data sets. Usually used by management to do queries on data (such as current performance vs targets etc)	

Amazon Data Warehousing (REDSHIFT)

- Databases use different type of architecture both from a database perspective and infrastructure layer
- Amazon's Data Warehouse Solution is called REDSHIFT

ElastiCache (Caches most common queries)

- Is a web service that makes it easy to deploy , operate, and scale an in-memory cache in the cloud. The service improves the performance of web applications by allowing you to retrieve information from fast, managed, in-memory caches, instead of relying entirely on slower disk-based databases.
- ElastiCache supports two open-source in-memory caching engines:
- ElastiCache supports 2 open-source in-memory caching engines
 - Memcached
 - Redis
- Caching engine in the cloud for caching most common queries

TIPS

- RDS (SQL/OLTP)
 - SQL
 - MySQL
 - PostgreSQL
 - Oracle
 - Aurora
 - Amazon Proprietary DB 6 copies of Data across 3 availability zones
 - MariaDB
- AMAZON NoSQL Solution: DynamoDB (No SQL)
- Redshift OLAP Amazon's Data Warehousing solution, usecase online analytics processing, BI tools, Making really complex queries.
- ElastiCache
 - Memcached

- Redis
- Is an in-memory caching engine in the cloud and take load of production DB.
- ElastiCache to speed up performance of existing database (Frequent identical queries)

Databases 101:

- Provision an RDS Instance
- Open MySQL Port to Web-DMZ SG
 - Allows our web server to talk to our DB VIA the port 3306 (mysql port)
 - Opening up a firewall, allowing your web server to talk to your DB
- Create and EC2 Instance
- Install Wordpress Using Boot Strap Script
 - Is going to automate entire installation of WP
- Registrar the EC2 Instance to the Target Group
 - Allowing to put our EC2 behind an application load balancer
 - Load balancer
 - Instead of having and IP address for your EC2 instance, you get a DNS name,
 - DNS name maps the application load balancer back to your EC2 instance
- Updated Wordpress to the DNS name of ALB
- Take a Snapshot
 - To create auto scaling group which will scale out automatically based on load, and other requirements we set, that will allow us to have a fault tolerant WP site