Answer Sheet

Q1)

Your company is migrating infrastructure to AWS. A large number of developers and administrators will need to control this infrastructure using the AWS Management Console.

The Identity Management team is objecting to creating an entirely new directory of IAM users for all employees, and the employees are reluctant to commit yet another password to memory.

Which of the following will satisfy both these stakeholders?

- Users sign in using an OpenID Connect (OIDC) compatible IdP, receive an authentication token, then use that token to log in to the AWS Management Console.
- Users log in to the AWS Management Console using the AWS Command Line Interface.
- ✅ Users request a SAML assertion from your on-premises SAML 2.0-compliant identity provider (IdP) and use that assertion to obtain federated access to the AWS Management Console via the AWS single sign-on (SSO) endpoint.

Explanation:-it uses the on-premises 2.0 SAML compliant IdP and get the federated access to the AWS, thus avoiding creating any IAM User for the employees in the organization.

- Users log in directly to the AWS Management Console using the credentials from your on-premises Kerberos compliant Identity provider.

Q2)

You are managing the AWS account of a big organization. The organization has more than 1000+ employees and they want to provide access to the various services to most of the employees.

Which of the below-mentioned options is the best possible solution in this case?

- The user should create a separate IAM user for each employee and provide access to them as per the policy.
- The user should create an IAM role and attach STS with the role. The user should attach that role to the EC2 instance and setup AWS authentication on that server.
- The user should create IAM groups as per the organization's departments and add each user to the group for better access control.
- ✅ Attach an IAM role with the organization's authentication service to authorize each user for various AWS services.

Explanation:-The best practice for IAM is to create roles which have specific access to an AWS service and then give the user permission to the AWS service via the role. it authenticates the users with the organization's authentication service and creates an appropriate IAM Role for accessing the AWS services. For the best practices on IAM policies

---

Q3)

A company host a web application. There is a requirement that users from the internet must be able to access this application, but the application server itself should not be exposed to the internet.

Which of the below options can help to fulfill these requirements?

- Configure an IPsec VPN connection, and provide the users with the configuration details. Create a public subnet in your VPC, and place your application servers in it.
- Implement Elastic Load Balancing with an SSL listener that terminates the back-end connection to the application.
- Implement AWS Direct Connect, and create a private interface to your VPC. Create a public subnet and place your application servers in it.
- ✅ Configure an SSL VPN solution in a public subnet of your VPC, then configure VPN client software on all user computers. Create a private subnet in your VPC and place your application servers in it.

Explanation:-configuring the SSL VPN solution is cost-effective and allows access only to the business travelers and since the application servers are in private subnet, the application is not accessible via the internet.

Q4)

A customer is running an application in the US-West region and wants to set up disaster recovery failover to Singapore region.

The customer is interested in achieving a low RPO for an RDS multi-AZ DB instance.

Which approach is best suited to this need?

- Copying of RDS incremental snapshots
- ✅ Asynchronous replication

Explanation:-When you have cross-region replication for RDS, this is done Asynchronously. Having Synchronous replication would be too much of an overhead for a cross-region replication.

- Synchronous replication
- Route53 health checks

Q5)

A newspaper organization has a requirement to store around 20TB of data for their readers. This data compromises of newspapers in various languages.

They wanted to use a search feature for users to search for articles on the site.

Which AWS service can help to fulfill this requirement?

- CloudFront
- ✅ CloudSearch

Explanation:-With Amazon CloudSearch, you can quickly add rich search capabilities to your website or application. You don't need to become a search expert or worry about hardware provisioning, setup, and maintenance. With a few clicks in the AWS Management Console, you can create a

search domain and upload the data that you want to make searchable, and Amazon CloudSearch will automatically provision the required resources and deploy a highly tuned search index. You can easily change your search parameters, fine tune search relevance, and apply new settings at any time. As your volume of data and traffic fluctuates, Amazon CloudSearch seamlessly scales to meet your needs.

- Multi-AZ RDS
- Kinesis

Q6)

Due to a lot of your EC2 services going offline at least once a week for no apparent reason your security officer has told you that you need to tighten up the logging of all events that occur on your AWS account.

He wants to be able to access all events that occur on the account across all regions quickly and in the simplest way possible.

He also wants to make sure he is the only person that has access to these events in the most secure way possible.

Which of the following would be the best solution to assure his requirements are met? Choose the correct answer from the below options:

- Use CloudTrail to log all events to a separate S3 bucket in each region as CloudTrail cannot write to a bucket in a different region. Use MFA and bucket policies on all the different buckets.
- Use CloudTrail to log all events to an Amazon Glacier Vault. Make sure the vault access policy only grants access to the security officer's IP address.
- Use CloudTrail to send all API calls to CloudWatch and send an email to the security officer every time an API call is made. Make sure the emails are encrypted.
- ✅ Use CloudTrail to log all events to one S3 bucket. Make this S3 bucket only accessible by your security officer with a bucket policy that restricts access to his user only and also add MFA to the policy for a further level of security.

Explanation:-The main points to consider in this scenario is: (1) the security officer needs to access all events that occur on the account across all the regions, and (2) only that security officer should have the access. it configures only one S3 bucket for all the CloudTrail log events on the account across all the regions. It also restricts the access to the security officer only via the bucket policy.

Q7)

Someone on your team configured a Virtual Private Cloud with two public subnets in two separate AZs and two private subnets in two separate AZs. Each public subnet AZ has a matching private subnet AZ. The VPC and its subnets are properly configured.

You also notice that there are multiple webserver instances in the private subnet, and you've been charged with setting up a public-facing Elastic Load Balancer which will accept requests from clients and distribute those requests to the webserver instances.

How can you set this up without making any significant architectural changes? Choose the correct option from the below:

- You can't. Webserver instances must be in public subnets in order for this to work.
- Put the webserver instances in the public subnets and then configure the ELB with those subnets.
- ✅ Select both of the public subnets when configuring the ELB.

Explanation:-you need to associate the public subnets with the internet facing load balancer. You would also need to ensure that the security group that is assigned to the load balancer has the listener ports open and the security groups of the private instances allow traffic on the listener ports and the health check ports.

- Select both of the private subnets which contain the webserver instances when configuring the ELB.

Q8)

An AWS customer is deploying a web application that is composed of a front end running on Amazon EC2 and confidential data that is stored on Amazon S3. The customer's Security policy requires that the all-access operations to this sensitive data must be authenticated and authorized by a centralized access management system that is operated by a separate security team. In addition, the web application team that owns and administers the EC2 web front-end instances is prohibited from having any ability to access the data that circumvents this centralized access management system.

Which of the following configurations will support these requirements:

- ✅ Configure the web application to authenticate end users against the centralized access management system. Have the web application provision trusted users STS tokens entitling the download of approved data directly from Amazon S3.

Explanation:-the access to the sensitive data on Amazon S3 is only given to the authenticated users.

- Encrypt the data on Amazon S3 using a CloudHSM that is operated by the separate security team. Configure the web application to integrate with the CloudHSM for decrypting approved data access operations for trusted end users.
- Configure the web application to authenticate end users against the centralized access management system using SAML. Have the end users authenticate to IAM using their SAML token and download the approved data directly from Amazon S3.
- Have the separate security team create an IAM Role that is entitled to access the data on Amazon S3. Have the web application team provision their instances with this Role while denying their IAM users access to the data on Amazon S3.

Q9)

An online gaming server in which you have recently increased it's IOPS performance, by creating a RAID 0 configuration has now started to have bottleneck problems due to your instance bandwidth.

Which of the following would be the best solution for this to increase throughput? Choose the correct answer from the below options:

- Move all your EC2 instances to the same availability zone.
- Use instance store backed instances and stripe the attached ephemeral storage devices and use DRBD Asynchronous Replication.
- ✅ Use Single Root I/O Virtualization (SR-IOV) on all the instances.

Explanation:-SR-IOV helps in achieving higher network throughput, lower CPU utilization, and lower network latency which can translate into supporting more VMs per host, delivering increased network bandwidth utilization on the host, and providing greater performance predictability to the instances.

- Use a RAID 1 configuration instead of RAID 0.

Q10)

A company has an application that is hosted on an EC2 instance.

The code is written in .NET and connects to a MySQL RDS database.

If you're executing .NET code against AWS on an EC2 instance that is assigned an IAM role, which of the following is a true statement? Choose the correct option from the below:

- None of these
- ✅ The code will assume the same permissions as the EC2 role

Explanation:-The best practice for IAM is to create roles which have specific access to an AWS service and then give the user permission to the AWS service via the role. To get the role in place, follow the below steps Step 1) Create a role which has the required ELB access Step 2) You need to provide permissions to the underlying EC2 instances in the Elastic Load Balancer

- The code must have AWS access keys in order to execute
- Only .NET code can assume IAM roles

Q11)

You tried to integrate 2 systems (front end and back end) with an HTTP interface to one large system. These subsystems don't store any state inside. All state is stored in a DynamoDB table.

You have launched each of the subsystems with separate AMIs.

After testing, these servers stopped running and are issuing malformed requests that do meet the HTTP specifications of the client.

Your developers fix the issue and deploy the fix to the subsystems as soon as possible without service disruption.

What are the 3 most effective options from the below to deploy these fixes?

- Use Amazon Cloudfront with access the front end server with origin fetch.
- ✅ Use Elastic Load balancing in front of the back-end system and Auto scaling to keep the specified number of instances

Explanation:-you can pause instances in AutoScaling, apply the patches and then add the instances back to AutoScaling and it will be registered with ELB.

- ✅ Use Elastic Load balancing in front of the front-end system and Auto scaling to keep the specified number of instances

Explanation:-you can pause instances in AutoScaling, apply the patches and then add the instances back to AutoScaling and it will be registered with ELB.

- ✅ Use AWS Opsworks autohealing for both the front end and back end instance pair

Explanation:-Autohealing would try to bring the instances back up with the patches deployed. Please see the "More information.." section.

- Use VPC
- Use Amazon SQS between the front end and back end subsystems.

Q12)

You are maintaining an application that is spread across multiple web servers and has incoming traffic balanced by ELB. The application allows users to upload pictures. Currently, each web server stores the image and a background task synchronizes the data between servers.

However, the synchronization task can no longer keep up with the number of images uploaded What change could you make so that all web servers have a place to store and read images at the same time? Choose an option from the below:

- Store the images on Amazon EBS
- ✅ Store the images in Amazon S3

Explanation:-S3 provides a durable, secure, cost effective, and highly available storage service for the uploaded pictures.

- Store the images on the ELB
- Store the images on Amazon Cloudfront

Q13)

As a solution architect professional, you have been requested to launch 20 Large EC2 instances which will all be used to process huge amounts of data.

There is also a requirement that these instances will need transfer data back and forth between each other.

Which of the following would be the most efficient setup to achieve this? Choose the correct option from the below:

- Ensure all instances are placed in the same Availability Zone.
- Use the largest EC2 instances currently available on AWS and make sure they are spread across multiple availability zones.
- ✅ Use Placement Groups and ensure all instances are launched at the same time.

Explanation:-Placement Group enables applications to get the low-latency network performance necessary for tightly-coupled node-to-node communication typical of many high-performance computing applications.

- Ensure all the instances are placed in the same region.

Q14)

A company has placed a set of on-premise resources with an AWS Direct Connect provider. After establishing connections to a local AWS region in the US, the company needs to establish a low latency dedicated connection to an S3 public endpoint over the Direct Connect dedicated low latency connection.

What steps need to be taken to accomplish configuring a direct connection to a public S3 endpoint? Choose the correct answer from the options given below:

Add a BGP route as part of the on-premise router; this will route S3 related traffic to the public S3 endpoint to dedicated AWS region.

Configure a public virtual interface to connect to a public S3 endpoint resource.

Explanation:-You can create a public virtual interface to connect to public resources or a private virtual interface to connect to your VPC. You can configure multiple virtual interfaces on a single AWS Direct Connect connection, and you'll need one private virtual interface for each VPC to connect to. Each virtual interface needs a VLAN ID, interface IP address, ASN, and BGP key. connect to S3 endpoint, a public virtual interface needs to be created, not private.

Establish a VPN connection from the VPC to the public S3 endpoint.

Configure a private virtual interface to connect to the public S3 endpoint via the Direct Connect connection.

Q15)

You have been asked to manage your AWS infrastructure In a manner similar to the way you might manage application code.

You want to be able to deploy exact copies of different versions of your infrastructure, stage changes into different environments, revert back to previous versions, and identify what versions are running at any particular time (development test QA . production).

Which approach addresses this requirement?

Use cost allocation reports and AWS Opsworks to deploy and manage your infrastructure.

Use AWS CloudWatch metrics and alerts along with resource tagging to deploy and manage your infrastructure.

Use AWS CloudFormation and a version control system like GIT to deploy and manage your infrastructure.

Explanation:-You can use AWS Cloud Formation's sample templates or create your own templates to describe the AWS resources, and any associated dependencies or runtime parameters, required to run your application. You don't need to figure out the order for provisioning AWS services or the subtleties of making those dependencies work. CloudFormation takes care of this for you. After the AWS resources are deployed, you can modify and update them in a controlled and predictable way, in effect applying version control to your AWS infrastructure the same way you do with your software. You can also visualize your templates as diagrams and edit them using a drag-and-drop interface with the AWS CloudFormation Designer. AWS CloudFormation gives developers and systems administrators an easy way to create and manage a collection of related AWS resources, provisioning and updating them in an orderly and predictable fashion.

Use AWS Beanstalk and a version control system like GIT to deploy and manage your infrastructure.

Q16)

While hosting a static website with Amazon S3, your static JavaScript code attempts to include resources from another S3 bucket but permission is denied.

How might you solve the problem? Choose the correct option from the below:

Disable Public Object Permissions

Move the object to the main bucket

None of these

Enable CORS Configuration

Explanation:-the instance in the EU region will not have any changes made after copying the AMI. You will need to copy the AMI#2 to eu-west-1 and then launch the instance again to have all the changes.

Q17)

A user has launched an EC2 instance store-backed instance in the us-east-1a zone. The user created AMI #1 and copied it to the eu-west-1 region.

After that, the user made a few updates to the application running in the us-east-1a zone.

The user makes an AMI #2 after the changes.

If the user launches a new instance in Europe from the AMI #1 copy, which of the below-mentioned statements is true?

The new instance will have the changes made after the AMI copy since AWS keeps updating the AMI.

The new instance in the eu-west-1 region will not have the changes made after the AMI copy.

Explanation:-the instance in the EU region will not have any changes made after copying the AMI. You will need to copy the AMI#2 to eu-west-1 and then launch the instance again to have all the changes.

The new instance will have the changes made after the AMI copy as AWS just copies the reference of the original AMI during the copying. Thus, the copied AMI will have all the updated data.

It is not possible to copy the instance store backed AMI from one region to another.

Q18)

You've been tasked with moving an e-commerce web application from a customer's data center into a VPC. The application must be fault tolerant and well as highly scalable.

Moreover, the customer is adamant that service interruptions not affect the user experience.

As you near launch, you discover that the application currently uses multicast to share session state between web servers.

In order to handle session state within the VPC, you choose to which of the following option:

Create a mesh VPN between instances and allow multicast on it.

Store session state in Amazon ElastiCache for Redis.

Explanation:-Mesh VPN is just not fault tolerant or highly scalable - the client's real priorities. It's failure would impact users. The supernode that handles the registration is a single point of failure and in case of failure, new VPN nodes would not be able to register. Also, the nodes would't register across multiple AZs. Even if it is possible it is very cumbersome.

Store session state in Amazon Relational Database Service.

Enable session stickiness via Elastic Load Balancing.

Q19)

You've created a temporary application that accepts image uploads, stores them in S3, and records information about the image in RDS. After building this architecture and accepting images for the duration required, it's time to delete the CloudFormation template.

However, your manager has informed you that for archival reasons the RDS data needs to be stored and the S3 bucket with the images needs to remain.

Your manager has also instructed you to ensure that the application can be restored by a CloudFormation template and run next year during the same period.

Knowing that when a CloudFormation template is deleted, it will remove the resources it created.

What is the best method for achieving the desired goals? Choose the correct option from the below:

Set the DeletionPolicy on the S3 resource to snapshot and the DeletionPolicy on the RDS resource to snapshot.
For both the RDS and S3 resource types on the CloudFormation template, set the DeletionPolicy to retain.
Enable S3 bucket replication on the source bucket to a destination bucket to maintain a copy of all the S3 objects, set the deletion policy for the RDS instance to snapshot.
✅ Set the DeletionPolicy on the S3 resource declaration in the CloudFormation template to retain, set the RDS resource declaration DeletionPolicy to snapshot.
Explanation:-The main points in this questions are: (i) need for an ability by which the RDS data that is stored and can be restored of needed and (ii) the S3 bucket with the images needs to retain. it uses retain policy for S3 bucket and snapshot policy for RDS such that the data can be restored when needed. More information on DeletionPolicy Options: Delete AWS CloudFormation deletes the resource and all its content if applicable during stack deletion. Retain AWS CloudFormation keeps the resource without deleting the resource or its contents when its stack is deleted. Snapshot For resources that support snapshots (AWS::EC2::Volume, AWS::ElastiCache::CacheCluster, AWS::ElastiCache::ReplicationGroup, AWS::RDS::DBInstance, AWS::RDS::DBCluster, and AWS::Redshift::Cluster), AWS CloudFormation creates a snapshot for the resource before deleting it. For more information on CloudFormation deletion policy,

---

Q20) Which of the following cache options is normally preferred when building gaming applications?

Primary Cache
None of these
Memcached
✅ Redis
Explanation:-ElastiCache for Redis is the best suited for gaming applications since leaderboards, sessions and profiles are the top functions for game developers, its raw event stream is used for dashboards and powering interactive customized campaigns, on top of being consumed by downstream processes for deeper analytics and long-term storage. Also, for sorting the results for top performance or scores, Redis' data structure is very helpful.

Q21)

One of your requirements is to setup an S3 bucket to store your files like documents and images.

However, those objects should not be directly accessible via the S3 URL, they should only be accessible from pages on your website so that only your paying customers can see them. How could you implement this? Choose the correct option from the below:

Use HTTPS endpoints to encrypt your data.
You can use server-side and client-side encryption, where only your application can decrypt the objects
✅ You can use a bucket policy and check for the AWS: Referer key in a condition, where that key matches your domain
Explanation:-Suppose you have a website with the domain name (www.example.com or example.com) with links to photos and videos stored in your S3 bucket, examplebucket. By default, all the S3 resources are private, so only the AWS account that created the resources can access them. To allow read access to these objects from your website, you can add a bucket policy that allows s3:GetObject permission with a condition, using theaws:referer key, that the get request must originate from specific web pages. it defines appropriate bucket policy to give the access to the S3 content to the authenticated users.
You can't. The S3 URL must be public in order to use it on your website.

Q22)

An Enterprise customer is starting their migration to the cloud, their main reason for migrating is agility, and they want to make their internal Microsoft Active Directory available to any applications running on AWS; this is so internal users only have to remember one set of credentials and as a central point of user control for leavers and joiners.

How could they make their Active Directory secure, and highly available, with minimal on-premises infrastructure changes, in the most cost and time-efficient way? Choose the most appropriate option from the below:

✅ Using VPC, they could create an extension to their data center and make use of resilient hardware IPSec tunnels; they could then have two domain controller instances that are joined to their existing domain and reside within different subnets, in different Availability Zones.
Explanation:-using an IPSec tunnel can help decrypt all the traffic from the on-premise to AWS. The domain controllers in separate AZ's can address high availability.
The customer could create a stand-alone VPC with its own Active Directory Domain Controllers; two domain controller instances could be configured, one in each Availability Zone; new applications would authenticate with those domain controllers.
Using Amazon Elastic Compute Cloud (EC2), they could create a DMZ using a security group; within the security group they could provision two smaller Amazon EC2 instances that are running Openswan for resilient IPSec tunnels, and two larger instances that are domain controllers; they would use multiple Availability Zones.
Within the customer's existing infrastructure, they could provision new hardware to run Active Directory Federation Services; this would present Active Directory as a SAML2 endpoint on the internet; any new application on AWS could be written to authenticate using SAML2

Q23)

Currently, a company uses Redshift to store its analyzed data. They have started with the base configuration.

What would they get when they initially start using Redshift?

- Two nodes with 320GB each
- ✅ One node of 160GB
- Two nodes with 160GB each
- One node of 320GB

Q24) What can be done if a company wants to establish a low latency dedicated connection to an S3 public endpoint over the Direct Connect?

- Configure a private virtual interface to connect to the public S3 endpoint via the Direct Connect connection.
- Establish a VPN connection from the VPC to the public S3 endpoint.
- ✅ Configure a public virtual interface to connect to a public S3 endpoint resource.

Explanation:-You can create a public virtual interface to connect to public resources or a private virtual interface to connect to your VPC. You can configure multiple virtual interfaces on a single AWS Direct Connect connection, and you'll need one private virtual interface for each VPC to connect to. Each virtual interface needs a VLAN ID, interface IP address, ASN, and BGP key. it creates a public virtual interface to connect to S3 endpoint.

- Add a BGP route as part of the on-premise router; this will route S3 related traffic to the public S3 endpoint to dedicated AWS region.

Q25) Which of the following are some of the best examples where Amazon Kinesis can be used?

- Real-time data analytics
- Accelerated log and data feed intake
- ✅ All of these

Explanation:-The following are typical scenarios for using Kinesis Data Streams: Accelerated log and data feed intake and processing You can have producers push data directly into a stream. For example, push system and application logs and they are available for processing in seconds. This prevents the log data from being lost if the front end or application server fails. Kinesis Data Streams provides accelerated data feed intake because you don't batch the data on the servers before you submit it for intake. Real-time metrics and reporting You can use data collected into Kinesis Data Streams for simple data analysis and reporting in real time. For example, your data-processing application can work on metrics and reporting for system and application logs as the data is streaming in, rather than wait to receive batches of data. Real-time data analytics This combines the power of parallel processing with the value of real-time data. For example, process website clickstreams in real time, and then analyze site usability engagement using multiple different Kinesis Data Streams applications running in parallel. Complex stream processing You can create Directed Acyclic Graphs (DAGs) of Amazon Kinesis Data Streams applications and data streams. This typically involves putting data from multiple Amazon Kinesis Data Streams applications into another stream for downstream processing by a different Amazon Kinesis Data Streams application.

- Real-time metrics and reporting

Q26)

An organization is planning to use AWS for their production roll out. The organization wants to implement automation for deployment such that it will automatically create a LAMP stack, download the latest PHP installable from S and setup the ELB.

Which of the below mentioned AWS services meets the requirement for making an orderly deployment of the software?

- AWS Cloudfront
- AWS Cloudformation
- AWS DevOps
- ✅ AWS Elastic Beanstalk

Explanation:-The Elastic Beanstalk is an easy-to-use service for deploying and scaling web applications and services. We can simply upload code and Elastic Beanstalk automatically handles the deployment, from capacity provisioning, load balancing, auto-scaling to application health monitoring. Meanwhile we can retain full control over the AWS resources used in the application and can access the underlying resources at any time.

Q27) An organization has created one IAM user and applied the below-mentioned policy to the user. What entitlements do the IAM users avail with this policy? { “Version”: “2012-10-17”, “Statement”: [{ “Effect”: “Allow”, “Action”: “ec2:Describe*”, “Resource”: “*” }, { “Effect”: “Allow” “Action”: [“cloudwatch:ListMetrics”, “cloudwatch:GetMetricStatistics”, “cloudwatch:Describe*”], “Resource”: “*” }, { “Effect”: “Allow”, “Action”: “autoscaling:Describe*”, “Resource”: “*” }] }

- The policy will allow the user to perform all read-only activities on the EC2 services.
- The policy will allow the user to perform all read and write activities on the EC2 services.
- ✅ The policy will allow the user to perform all read-only activities on the EC2 services except load Balancing.

Explanation:-The above policy will allow the user to view EC2 instances, look at Autoscaling and Cloudwatch but not allow the user access to Load Balancing. For access to load balancing, you need to have the following statements as well. "Effect": "Allow", "Action": "elasticloadbalancing:Describe*", "Resource": "*"

- The policy will allow the user to list all the EC2 resources except EBS.

Q28)

You have an EBS root device on /dev/sda1 on one of your EC2 instances. You are having trouble with this particular instance and you want to either Stop/Start, Reboot or Terminate the instance but you do not want to lose any data that you have stored on /dev/sda1.

Which of the below statements best describes the effect each change of instance state would have on the data you have stored on /dev/sda1? Choose the correct option from the below:

- Whether you stop/start, reboot or terminate the instance it does not matter because data on an EBS volume is not ephemeral and the data will

not be lost regardless of what method is used

Whether you stop/start, reboot or terminate the instance it does not matter because data on an EBS volume is ephemeral and it will be lost no matter what method is used.

If you stop/start the instance the data will not be lost. However, if you either terminate or reboot the instance the data will be lost.

The data in an instance store is not permanent - it persists only during the lifetime of the instance. The data will be lost if you terminate the instance. However, the data will remain on /dev/sda1 if you reboot or stop/start the instance because data on an EBS volume is not ephemeral.

Explanation:-Since this is an EBS backed instance, it can be stopped and later restarted without affecting data stored in the attached volumes. By default, the root device volume for this intance will b deleted when the instance terminates.

Q29)

Your company hosts an on-premises legacy engineering application with 900GB of data shared via a central file server. The engineering data consists of thousands of individual files ranging in size from megabytes to multiple gigabytes.

Engineers typically modify 5-10 percent of the files a day.

Your CTO would like to migrate this application to AWS, but only if the application can be migrated over the weekend to minimize user downtime.

You calculate that it will take a minimum of 48 hours to transfer 900GB of data using your company's existing 45-Mbps Internet connection.

After replicating the application's environment in AWS, which option will allow you to move the application's data to AWS without losing any data and within the given timeframe?

Copy the data to Amazon S3 using multiple threads and multi-part upload for large files over the weekend, and work in parallel with your developers to reconfigure the replicated application environment to leverage Amazon S3 to serve the engineering files.

Leverage the AWS Storage Gateway to create a Gateway-Stored volume. On Friday copy the application data to the Storage Gateway volume. After the data has been copied, perform a snapshot of the volume and restore the volume as an EBS volume to be attached to your AWS file server on Sunday.

Sync the application data to Amazon S3 starting a week before the migration, on Friday morning perform a final sync, and copy the entire data set to your AWS file server after the sync completes.

Explanation:-In this scenario, following important points need to be considered - (i) only fraction of the data (5-10%) is modified every day, (ii) there are only 48 hrs for the migration, (iii) downtime should be minimized, and (iv) there should be no data loss. the data changes are limited and can be propagated over the week. Also, the bandwidth would be used efficiently, and you would have sufficient time and bandwidth in hand, should there be any unexpected issues while migrating.

Copy the application data to a 1-TB USB drive on Friday and immediately send overnight, with Saturday delivery, the USB drive to AWS Import/Export to be imported as an EBS volume, mount the resulting EBS volume to your AWS file server on Sunday

---

Q30)

An application is deployed in multiple Availability Zones in a single region. In the event of failure, the RTO must be less than 3 hours, and the RPO is 15 minutes.

Which DR strategy can be used to achieve this RTO and RPO in the event of this kind of failure?

Use synchronous database master-slave replication between two Availability Zones

Take hourly DB backups to Amazon S3, with transaction logs stored in S3 every 5 minutes

Explanation:-it takes hourly backups to Amazon S3 - which makes restoring the backups quick, and since the transaction logs are stored in S3 every 5 minutes, it will help to restore the application to a state that is within the RPO of 15 minutes.

Take hourly DB backups to an Amazon EC2 instance store volume, with transaction logs stored in Amazon S3 every 5 minutes.

Take 15-minute DB backups stored in Amazon Glacier, with transaction logs stored in Amazon S3 every 5 minutes

---

Q31)

There is a requirement to move a legacy app to AWS. This legacy app still requires accessing some of the services on the on-premise architecture.

Which of the 2 below options will help to fulfill this requirement?

An AWS Direct Connect link between the VPC and the network housing the internal services.

Explanation:-The scenario requires you to connect your on-premise server/instance with Amazon VPC. When such scenarios are presented, always think about services such as Direct Connect, VPN, and VM Import and Export as they help either connecting the instances from different location or importing them from one location to another. Direct Connect sets up a dedicated connection between on-premise data-center and Amazon VPC, and provides you with the ability to connect your on-premise servers with the instances in your VPC.

An Internet Gateway to allow a VPN connection.

An Elastic IP address on the VPC instance.

An IP address space that does not conflict with the one on-premises

Explanation:-The scenario requires you to connect your on-premise server/instance with Amazon VPC. When such scenarios are presented, always think about services such as Direct Connect, VPN, and VM Import and Export as they help either connecting the instances from different location or importing them from one location to another. there should not be a conflict between IP address of on-premise servers and the instances in VPC for them to communicate.

Q32)

Your customer is implementing a video-on-demand streaming platform on AWS. The requirement is to be able to support multiple devices such as iOS, Android, and Windows as client devices, using a standard client player, using streaming technology and scalable architecture with cost-effectiveness.

Which architecture meets the requirements?

Launch a streaming server on Amazon Elastic Compute Cloud (EC2) (for example, Adobe Media Server), and store the video contents as an

origin server. Configure the Amazon CloudFront distribution with a download option to steam the video contents.
- Launch a steaming server on Amazon EC2 (for example, Adobe Media Server), and store the video contents as an origin server. Launch and configure the required amount of streaming servers on Amazon EC2 as an edge server to stream the video contents.
- Store the video contents to Amazon Simple Storage Service (S3) as an origin server. Configure the Amazon CloudFront distribution with a streaming option to stream the video contents.
- ✅ Store the video contents to Amazon S3 as an origin server. Configure the Amazon CloudFront distribution with a download option to stream the video contents

Explanation:-(a) it uses CloudFront distribution with download option for streaming the on demand videos using HLS on any mobile, and (b) it uses S3 as origin, so keeps the cost low.

Q33)

Company B has created an e-commerce site using DynamoDB and is designing a table named Products that includes items purchased and the users who purchased them.

When creating a primary key on this table which of the following would be the best attribute? Select the best possible answer:

- product_id where there are few products to many users
- ✅ user_id where there are many users to few products

Explanation:-When defining primary keys, you should always use the "many to few principle".
- category_id where there are few categories to many products
- None of these

Q34)

You have multiple instances behind private and public subnets.

None of the instances have an EIP assigned to them.

How can you securely connect them to the internet just to be able to download system updates? Choose the correct option from the below:

- Connect to a VPN
- Assign EIP to each instance
- ✅ Create a NAT instance

Explanation:-You can use a Network Address Translation (NAT) instance in a public subnet in your VPC to enable instances in the private subnet to initiate outbound IPv4 traffic to the Internet or other AWS services, but prevent the instances from receiving inbound traffic initiated by someone on the Internet. All the instances in the private (and public) can get the system updates via NAT instance which is placed in the public subnet, without exposing themselves over the internet.
- Use both a NAT instance and a VPN

Q35)

A user is using Cloudformation to launch an EC2 instance and then configure an application after the instance is launched.

The user wants the stack creation of ELB and AutoScaling to wait until the EC2 instance is launched and configured properly.

How can the user configure this?

- The user can use the DependentCondition resource to hold the creation of the other dependent resources.
- The user can use the HoldCondition resource to wait for the creation of the other dependent resources.
- ✅ The user can use the WaitCondition resource to hold the creation of the other dependent resources.

Explanation:-You can use a wait condition for situations like the following: To coordinate stack resource creation with configuration actions that are external to the stack creation To track the status of a configuration process
- It is not possible that the stack creation will wait until one service is created and launched.

Q36)

A user has configured an SSL listener at ELB as well as on the back-end instances.

Which of the below-mentioned statements helps the user understand ELB traffic handling with respect to the SSL listener?

- ELB will intercept the request to add the cookie details if sticky session is enabled.
- It is not possible to have the SSL listener both at ELB and back-end instances.
- ✅ ELB will not modify the headers.

Explanation:-As per the AWS documentation, please find the below excerpts:
- ELB will modify headers to add requestor details.

Q37)

An Amazon Redshift cluster with four nodes is running 24/7/365 and expects potentially to add one on-demand node for one to two days once during the year.

Which architecture would have the lowest possible cost for the cluster requirement? Choose the correct answer from the below options:

- Purchase 5 reserved nodes to cover all possible usage during the year.
- Purchase 2 reserved nodes and utilize 3 on-demand nodes only for peak usage times.
- ✅ Purchase 4 reserved nodes and rely on on-demand instances for the fifth node, if required.

Explanation:-(a) the application requires 4 nodes throughout the year and reserved instances would save the cost, and (b) since the need of the other node is not assured, on-demand instance(s) can be purchased if and when needed.

Purchase 4 reserved nodes and bid on spot instances for the extra node if required.

Q38)

You decide to create a bucket on AWS S3 called 'mybucket' and then perform the following actions in the order that they are listed here.

- You upload a file to the bucket called 'file1'

- You upload a file called 'file2'

- You upload a file called 'file3'

- You upload another file called 'file2'

Which of the following is true for 'mybucket'? Choose the correct option from the below:

There will be 1 version ID for file1, there will be 2 version IDs for file2 and 1 version ID for file3
There will be 1 version ID for file1, the version ID for file2 will be null and there will be 1 version ID for file3
All file version ID's will be null because versioning must be enabled before uploading objects to 'mybucket'
✅ The version ID for file1 will be null, there will be 2 version IDs for file2 and 1 version ID for file3

Explanation:-Objects stored in your bucket before you set the versioning state have a version ID of null. When you enable versioning, existing objects in your bucket do not change. What changes is how Amazon S3 handles the objects in future requests. the file1 was put in the bucket before the versioning was enabled; hence, it will have null version ID. The file2 will have two version IDs, and file3 will have a single version ID.

Q39)

A corporate web application is deployed within an Amazon VPC and is connected to the corporate data center via IPSec VPN.

The application must authenticate against the on-premise LDAP server.

Once authenticated, logged-in users can only access an S3 keyspace specific to the user. Choose 2 options from the below:

✅ Develop an identity broker that authenticates against LDAP and then calls IAM Security Token Service (STS) to get IAM federated user credentials. The application calls the identity broker to get IAM federated user credentials with access to the appropriate S3 bucket.

Explanation:-There are two architectural considerations here: (1) The users must be authenticated via the on-premise LDAP server, and (2) each user should have access to S3 only. With this information, it is important to first authenticate the users using LDAP, get the IAM Role name, then get the temporary credentials from STS, and finally access the S3 bucket using those credentials. And second, create an IAM Role that provides access to S3. it follows the correct sequence. It develops an identity broker that authenticates users against LDAP, gets the security token from STS, and then accesses the S3 bucket using the IAM federated user credentials.

The application authenticates against LDAP the application then calls the AWS identity and Access Management (IAM) Security service to log in to IAM using the LDAP credentials the application can use the IAM temporary credentials to access the appropriate S3 bucket.
Develop an identity broker that authenticates against IAM Security Token Service (STS) to assume a IAM role in order to get temporary AWS security credentials The application calls the identity broker to get AWS temporary security credentials with access to the appropriate S3 bucket.
✅ The application authenticates against LDAP and retrieves the name of an IAM role associated with the user. The application then calls the IAM Security Token Service (STS) to assume that IAM role. The application can use the temporary credentials to access the appropriate S3 bucket.

Explanation:-There are two architectural considerations here: (1) The users must be authenticated via the on-premise LDAP server, and (2) each user should have access to S3 only. With this information, it is important to first authenticate the users using LDAP, get the IAM Role name, then get the temporary credentials from STS, and finally access the S3 bucket using those credentials. And second, create an IAM Role that provides access to S3. it follows the correct sequence. It authenticates users using LDAP, gets the security token from STS, and then accesses the S3 bucket using the temporary credentials.

Q40)

A company has hired you to assist with the migration of an interactive website that allows registered users to rate local restaurants. Updates to the ratings are displayed on the home page and ratings are updated in real time.

Although the website is not very popular today, the company anticipates that it will grow over the next few weeks.

They also want to ensure that the website to remain highly available.

The current architecture consists of a single Windows server 2008R2 web server and a MySQL database on Linux.

Both reside inside on an on-premise hypervisor.

What would be the most efficient way to transfer the application to AWS, ensuring high performance and availability?

Use AWS VM Import/Export to create an Amazon EC2 AMI of the web server. Configure auto-scaling to launch two web servers in us-west-1a and two in us-west-1b. Launch a Multi-AZ MySQL Amazon RDS instance in us-west-1b. Import the data into Amazon RDS from the latest MySQL backup. Use Amazon Route 53 to create a hosted zone and point an A record to the elastic load balancer.
Export web files to an Amazon 53 bucket in us-west-1. Run the website directly out of Amazon 53. Launch a multi-AZ MySQL Amazon RDS instance in us-west-1a. Import the data into Amazon RDS from the latest MySQL backup. Use Route 53 and create an alias record pointing to the elastic load balancer.
✅ Launch two Windows Server 2008 R2 instances in us-west-1b and two in us-west-1a and configure auto-scaling. Copy the web files from on premises web server to each Amazon EC2 web server, using Amazon S3 as the repository. Launch a multi -AZ MySQL Amazon RDS instance in us-west-1a. Import the data into Amazon RDS from the latest MySQL backup. Create an elastic load balancer (ELB) to front your web servers. Use Route 53 and create an alias record pointing to the ELB.

Explanation:-The main consideration in the question is that the architecture should be highly available with high performance. (a) EC2 servers can communicate with S3 for the web files, and (b) auto-scaling of web servers and the setup of Multi-AZ RDS instance as well as the Route 53 alias record with ELB provides high availability.
Use AWS VM Import/Export to create an Amazon EC2 AMI of the web server. Configure auto-scaling to launch two web servers in us-west-1a and two in us-west-1b. Launch a multi-AZ MySQL Amazon RDS instance in us-west-1a. Import the data Into Amazon RDS from the latest MySQL backup. Create an elastic load balancer (ELB) in front of your web servers. Use Amazon Route 53 and create an A record pointing to the ELB.

Q41)

A user has launched a large EBS backed EC2 instance in the US-East-1a region.

The user wants to achieve Disaster Recovery (DR. for that instance by creating another small instance in Europe.

How can the user achieve DR?

- Use the “Launch more like this” option to copy the instance from one region to another.
- ✅ Create an AMI of the instance and copy the AMI to the EU region. Then launch the instance from the EU AMI.

Explanation:-if you need an AMI across multiple regions, then you have to copy the AMI across regions. Note that by default AMI’s that you have created will not be available across all regions.

- Copy the running instance using the “Instance Copy” command to the EU region.
- Copy the instance from the US East region to the EU region.

Q42)

You are designing multi-region architecture and you want to send users to a geographic location based on latency-based routing, which seems simple enough; however, you also want to use weighted-based routing among resources within that region.

Which of the below setups would best accomplish this? Choose the correct answer from the below options:

- You will need to use AAAA - IPv6 addresses when you define your weighted based record sets.
- You will need to use complex routing (nested record sets) and ensure that you define the latency based records first
- ✅ You will need to use complex routing (nested record sets) and ensure that you define the weighted resource record sets first.

Explanation:-You must create the records in reverse order when creating complex routing. The discussion provided does not discuss this. It discusses how complex routing works. It doesn't talk about how to set it up. Please refer to the below documentation from AWS for an example where you can define complex routing

- This cannot be done. You can't use different routing records together.

Q43)

You are designing security inside your VPC. You are considering the options for establishing separate security zones, and enforcing network traffic rules across the different zones to limit which instances can communicate.

How would you accomplish these requirements? Choose 2 options from the below:

- Configure a security group for every zone. Configure a default allow all rule. Configure explicit deny rules for the zones that shouldn’t be able to communicate with one another.
- ✅ NACLs to explicitly allow or deny communication between the different IP address ranges, as required for inter zone communication.

Explanation:-you can explicitly allow or deny traffic based on certain IP address range.

- Configure multiple subnets in your VPC, one for each zone. Configure routing within your VPC in such a way that each subnet only has routes to other subnets with which it needs to communicate, and doesn’t have routes to subnets with which it shouldn’t be able to communicate.
- ✅ Configure a security group for every zone. Configure allow rules only between zones that need to be able to communicate with one another. Use the implicit deny all rule to block any other traffic.

Explanation:-Security Group in this case would act like a Firewall that provides security and control at the port/protocol level, and have "implicit deny all" rule and only allow what is needed.

Q44)

An auditor needs read-only access to all AWS resources and logs ofall the events that have occurred on AWS.

What is the best way for creating this sort of access? Choose the correct answer from the options below:

- Create an SNS notification that sends the CloudTrail log files to the auditor's email when CloudTrail delivers the logs to S3, but do not allow the auditor access to the AWS environment.
- The company should contact AWS as part of the shared responsibility model, and AWS will grant required access to the third-party auditor.
- Create a role that has the required permissions for the auditor.
- ✅ Enable CloudTrail logging and create an IAM user who has read-only permissions to the required AWS resources, including the bucket containing the CloudTrail logs.

Explanation:-sending the logs via email is not a good architecture.
