Amazon.Pre.AWS Certified Solutions Architect - Professional (SAP-C01).VCEup.105q - DEMO

Number: AWS
Passing Score: 800
Time Limit: 120 min

Certification: AWS Certified Solutions Architect - Professional
Certification Full Name: AWS Certified Solutions Architect - Professional
Certification Provider: Amazon

Exam A

QUESTION 1
Your company policies require encryption of sensitive data at rest. You are considering the possible options for protecting data while storing it at rest on an EBS data volume, attached to an EC2 instance.
Which of these options would allow you to encrypt your data at rest? (Choose 3)

A. Implement third party volume encryption tools
B. Implement SSL/TLS for all services running on the server
C. Encrypt data inside your applications before storing it on EBS
D. Encrypt data using native data encryption drivers at the file system level
E. Do nothing as EBS volumes are encrypted by default

Correct Answer: ACD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 2
A customer is deploying an SSL enabled web application to AWS and would like to implement a separation of roles between the EC2 service administrators that are entitled to login to instances as well as making API calls and the security officers who will maintain and have exclusive access to the application's X.509 certificate that contains the private key.

A. Upload the certificate on an S3 bucket owned by the security officers and accessible only by EC2 Role of the web servers.
B. Configure the web servers to retrieve the certificate upon boot from an CloudHSM is managed by the security officers.
C. Configure system permissions on the web servers to restrict access to the certificate only to the authority security officers
D. Configure IAM policies authorizing access to the certificate store only to the security officers and terminate SSL on an ELB.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
You'll terminate the SSL at ELB. and the web request will get unencrypted to the EC2 instance, even if the certs are stored in S3, it has to be configured on the web servers or load balancers somehow, which becomes difficult if the keys are stored in S3. However, keeping the keys in the cert store and using IAM to restrict access gives a clear separation of concern between security officers and developers. Developer's personnel can still configure SSL on ELB without actually handling the keys.

QUESTION 3
You have recently joined a startup company building sensors to measure street noise and air quality in urban areas. The company has been running a pilot deployment of around 100 sensors for 3 months each sensor uploads 1 KB of sensor data every minute to a backend hosted on AWS.
During the pilot, you measured a peak or 10 IOPS on the database, and you stored an =average of 3GB of sensor data per month in the database.
The current deployment consists of a load-balanced auto scaled Ingestion layer using EC2 instances and a PostgreSQL RDS database with 500GB standard storage.
The pilot is considered a success and your CEO has managed to get the attention or some potential investors. The business plan requires a deployment of at least 100K sensors which needs to be supported by the backend. You also need to store sensor data for at least two years to be able to compare year over year Improvements.
To secure funding, you have to make sure that the platform meets these requirements and leaves room for further scaling.
Which setup win meet the requirements?

A. Add an SQS queue to the ingestion layer to buffer writes to the RDS instance
B. Ingest data into a DynamoDB table and move old data to a Redshift cluster
C. Replace the RDS instance with a 6 node Redshift cluster with 96TB of storage
D. Keep the current architecture but upgrade RDS storage to 3TB and 10K provisioned IOPS

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 4
A web company is looking to implement an intrusion detection and prevention system into their deployed VPC. This platform should have the ability to scale to thousands of instances running inside of the VPC.
How should they architect their solution to achieve these goals?

A. Configure an instance with monitoring software and the elastic network interface (ENI) set to promiscuous mode packet sniffing to see an traffic across the VPC.
B. Create a second VPC and route all traffic from the primary application VPC through the second VPC where the scalable virtualized IDS/IPS platform resides.
C. Configure servers running in the VPC using the host-based 'route' commands to send all traffic through the platform to a scalable virtualized IDS/ IPS.
D. Configure each host with an agent that collects all network traffic and sends that traffic to the IDS/IPS platform for inspection.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 5
A company is storing data on Amazon Simple Storage Service (S3). The company's security policy mandates that data is encrypted at rest Which of the following methods can achieve this? (Choose 3)

A. Use Amazon S3 server-side encryption with AWS Key Management Service managed keys.
B. Use Amazon S3 server-side encryption with customer-provided keys.
C. Use Amazon S3 server-side encryption with EC2 key pair.
D. Use Amazon S3 bucket policies to restrict access to the data at rest.
E. Encrypt the data on the client-side before ingesting to Amazon S3 using their own master key.
F. Use SSL to encrypt the data while in transit to Amazon S3.

Correct Answer: ABE
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
http://docs.aws.amazon.com/AmazonS3/latest/dev/UsinaKMSEncrvption.html

QUESTION 6
Your firm has uploaded a large amount of aerial image data to S3. In the past, in your on-premises environment, you used a dedicated group of servers to oaten process this data and used Rabbit MQ - An open source messaging system to get job information to the servers. Once processed the data would go to tape and be shipped offsite. Your manager told you to stay with the current design, and leverage AWS archival storage and messaging services to minimize cost.
Which is correct?

A. Use SQS for passing job messages use Cloud Watch alarms to terminate EC2 worker instances when they become idle. Once data is processed, change the storage class of the S3 objects to Reduced Redundancy Storage.
B. Setup Auto-Scaled workers triggered by queue depth that use spot instances to process messages in SOS Once data is processed,
C. Change the storage class of the S3 objects to Reduced Redundancy Storage. Setup Auto-Scaled workers triggered by queue depth that use spot instances to process messages in SQS Once data is processed, change the storage class of the S3 objects to Glacier.
D. Use SNS to pass job messages use Cloud Watch alarms to terminate spot worker instances when they become idle. Once data is processed, change the storage class of the S3 object to Glacier.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 7
You've been hired to enhance the overall security posture for a very large e-commerce site. They have a well architected multi-tier application running in a VPC that uses ELBs in front of both the web and the app tier with static assets served directly from S3. They are using a combination of RDS and DynamoOB for their dynamic data and then archiving nightly into S3 for further processing with EMR. They are concerned because they found questionable log entries and suspect someone is attempting to gain unauthorized access.
Which approach provides a cost effective scalable mitigation to this kind of attack?

A. Recommend that they lease space at a DirectConnect partner location and establish a 1G DirectConnect connection to their VPC they would then establish Internet connectivity into their space, filter the traffic in hardware Web Application Firewall (WAF). And then pass the traffic through the DirectConnect connection into their application running in their VPC.
B. Add previously identified hostile source IPs as an explicit INBOUND DENY NACL to the web tier subnet.
C. Add a WAF tier by creating a new ELB and an AutoScaling group of EC2 Instances running a host-based WAF. They would redirect Route 53 to resolve to the new WAF tier ELB. The WAF tier would their pass the traffic to the current web tier The web tier Security Groups would be updated to only allow traffic from the WAF tier Security Group
D. Remove all but TLS 1.2 from the web tier ELB and enable Advanced Protocol Filtering. This will enable the ELB itself to perform WAF functionality.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 8
Your company is in the process of developing a next generation pet collar that collects biometric information to assist families with promoting healthy lifestyles for their pets Each collar will push 30kb of biometric data In JSON format every 2 seconds to a collection platform that will process and analyze the data providing health trending information back to the pet owners and veterinarians via a web portal Management has tasked you to architect the collection platform ensuring the following requirements are met.
Provide the ability for real-time analytics of the inbound biometric data
Ensure processing of the biometric data is highly durable. Elastic and parallel
The results of the analytic processing should be persisted for data mining
Which architecture outlined below win meet the initial requirements for the collection platform?

A. Utilize S3 to collect the inbound sensor data analyze the data from S3 with a daily scheduled Data Pipeline and save the results to a Redshift Cluster
B. Utilize Amazon Kinesis to collect the inbound sensor data, analyze the data with Kinesis clients and save the results to a Redshift cluster using EMR.
C. Utilize SQS to collect the inbound sensor data analyze the data from SQS with Amazon Kinesis and save the results to a Microsoft SQL Server RDS instance.
D. Utilize EM R to collect the inbound sensor data, analyze the data from EUR with Amazon Kinesis and save me results to DynamoDB.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 9
You are designing Internet connectivity for your VPC. The Web servers must be available on the Internet. The application must have a highly available architecture,
Which alternatives should you consider? (Choose 2)

A. Configure a NAT instance in your VPC Create a default route via the NAT instance and associate it with all subnets Configure a DNS A record that points to the NAT instance public IP address.
B. Configure a CloudFront distribution and configure the origin to point to the private IP addresses of your Web servers Configure a Route53 CNAME record to your CloudFront distribution.
C. Place all your web servers behind ELB Configure a Route53 CNMIE to point to the ELB DNS name.
D. Assign EIPs to all web servers. Configure a Route53 record set with all EIPs, with health checks and DNS failover.
E. Configure ELB with an EIP Place all your Web servers behind ELB Configure a Route53 A record that points to the EIP.

Correct Answer: CD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 10
Your team has a tomcat-based Java application you need to deploy into development, test and production environments. After some research, you opt to use Elastic Beanstalk due to its tight integration with your developer tools and RDS due to its ease of management. Your QA team lead points out that you need to roll a sanitized set of production data into your environment on a nightly basis. Similarly, other software teams in your org want access to that same restored data via their EC2 instances in your VPC.
The optimal setup for persistence and security that meets the above requirements would be the following.

A. Create your RDS instance as part of your Elastic Beanstalk definition and alter its security group to allow access to it from hosts in your application subnets.

B. Create your RDS instance separately and add its IP address to your application's DB connection strings in your code Alter its security group to allow access to it from hosts within your VPC's IP address block.
C. Create your RDS instance separately and pass its DNS name to your app's DB connection string as an environment variable. Create a security group for client machines and add it as a valid source for DB traffic to the security group of the RDS instance itself.
D. Create your RDS instance separately and pass its DNS name to your's DB connection string as an environment variable Alter its security group to allow access to It from hosts in your application subnets.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 11
Your company has an on-premises multi-tier PHP web application, which recently experienced downtime due to a large burst In web traffic due to a company announcement Over the coming days, you are expecting similar announcements to drive similar unpredictable bursts, and are looking to find ways to quickly improve your infrastructures ability to handle unexpected increases in traffic.
The application currently consists of 2 tiers a web tier which consists of a load balancer and several Linux Apache web servers as well as a database tier which hosts a Linux server hosting a MySQL database.
Which scenario below will provide full site functionality, while helping to improve the ability of your application in the short timeframe required?

A. Failover environment: Create an S3 bucket and configure it for website hosting. Migrate your DNS to Route53 using zone file import, and leverage Route53 DNS failover to failover to the S3 hosted website.
B. Hybrid environment: Create an AMI, which can be used to launch web servers in EC2. Create an Auto Scaling group, which uses the AMI to scale the web tier based on incoming traffic. Leverage Elastic Load Balancing to balance traffic between on-premises web servers and those hosted In AWS.
C. Offload traffic from on-premises environment: Setup a CloudFront distribution, and configure CloudFront to cache objects from a custom origin. Choose to customize your object cache behavior, and select a TTL that objects should exist in cache.
D. Migrate to AWS: Use VM Import/Export to quickly convert an on-premises web server to an AMI. Create an Auto Scaling group, which uses the imported AMI to scale the web tier based on incoming traffic. Create an RDS read replica and setup replication between the RDS instance and on¬premises MySQL server to migrate the database.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 12
You are implementing AWS Direct Connect. You intend to use AWS public service end points such as Amazon S3, across the AWS Direct Connect link. You want other Internet traffic to use your existing link to an Internet Service Provider.
What is the correct way to configure AWS Direct connect for access to services such as Amazon S3?

A. Configure a public Interface on your AWS Direct Connect link Configure a static route via your AWS Direct Connect link that points to Amazon S3 Advertise a default route to AWS using BGP.
B. Create a private interface on your AWS Direct Connect link. Configure a static route via your AWS Direct connect link that points to Amazon S3 Configure specific routes to your network in your VPC.
C. Create a public interface on your AWS Direct Connect link Redistribute BGP routes into your existing routing infrastructure; advertise specific routes for your network to AWS.
D. Create a private interface on your AWS Direct connect link. Redistribute BGP routes into your existing routing infrastructure and advertise a default route to AWS.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 13
Your application is using an ELB in front of an Auto Scaling group of web/application servers deployed across two AZs and a Multi-AZ RDS Instance for data persistence,
The database CPU is often above 80% usage and 90% of I/O operations on the database are reads. To improve performance you recently added a single¬node Memcached ElastiCache Cluster to cache frequent DB query results. In the next weeks the overall workload is expected to grow by 30%.
Do you need to change anything in the architecture to maintain the high availability or the application with the anticipated additional load? Why?

A. Yes, you should deploy two Memcached ElastiCache Clusters in different AZs because the RDS instance will not be able to handle the load if the cache node fails.
B. No, if the cache node fails you can always get the same data from the DB without having any availability impact.
C. No, if the cache node fails the automated ElastiCache node recovery feature will prevent any availability impact.
D. Yes, you should deploy the Memcached ElastiCache Cluster with two nodes in the same AZ as the RDS DB master instance to handle the load if one cache node fails.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
ElastiCache for Memcached
The primary goal of caching is typically to offload reads from your database or other primary data source. In most apps, you have hot spots of data that are regularly queried, but only updated periodically. Think of the front page of a blog or news site, or the top 100 leaderboard in an online game. In this type of case, your app can receive dozens, hundreds, or even thousands of requests for the same data before it's updated again. Having your caching layer handle these queries has several advantages. First, it's considerably cheaper to add an in-memory cache than to scale up to a larger database cluster. Second, an in-memory cache is also easier to scale out, because it's easier to distribute an in-memory cache
horizontally than a relational database.
Last, a caching layer provides a request buffer in the event of a sudden spike in usage. If your app or game ends up on
the front page of Reddit or the App Store, it's not unheard of to see a spike that is 10 to 100 times your normal
application load. Even if you autoscale your application instances, a 10x request spike will likely make your database very unhappy.
Let's focus on ElastiCache for Memcached first, because it is the best fit for a cachingfocused solution. We'll revisit Redis later in the paper, and weigh its advantages and disadvantages.
Architecture with ElastiCache for Memcached
When you deploy an ElastiCache Memcached cluster, it sits in your application as a separate tier alongside your
database. As mentioned previously, Amazon ElastiCache does not directly communicate with your database tier, or
indeed have any particular knowledge of your database. A simplified deployment for a web application looks something like this:

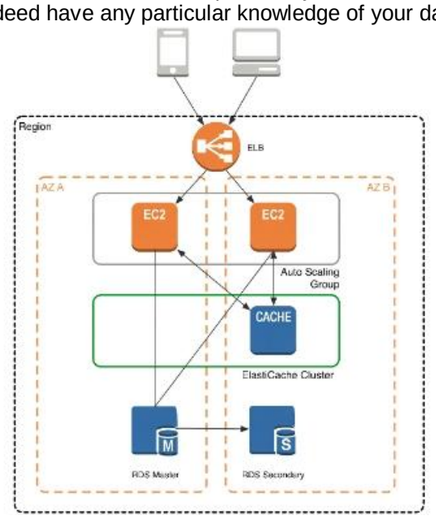

In this architecture diagram, the Amazon EC2 application instances are in an Auto Scaling group, located behind a load balancer using Elastic Load Balancing, which distributes requests among the instances. As requests come into a given EC2 instance, that EC2 instance is responsible for communicating with ElastiCache and the database tier. For development purposes, you can begin with a single ElastiCache node to test your application, and then scale to additional cluster nodes by modifying the ElastiCache cluster. As you add additional cache nodes, the EC2 application instances are able to distribute cache keys across multiple ElastiCache nodes. The most common practice is to use client- side sharding to distribute keys across cache nodes, which we will discuss later in this paper.

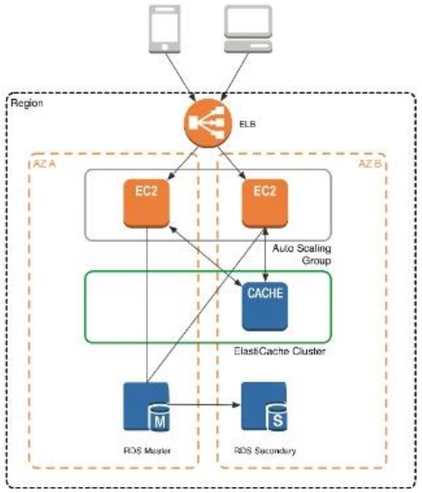

When you launch an ElastiCache cluster, you can choose the Availability Zone(s) that the cluster lives in. For best performance, you should configure your cluster to use the same Availability Zones as your application servers. To launch an ElastiCache cluster in a specific Availability Zone, make sure to specify the Preferred Zone(s) option during cache cluster creation. The Availability Zones that you specify will be where ElastiCache will launch your cache nodes. We recommend that you select Spread Nodes Across Zones, which tells ElastiCache to distribute cache nodes across these zones as evenly as possible. This distribution will mitigate the impact of an Availability Zone disruption on your ElastiCache nodes. The trade-off is that some of the requests from your application to ElastiCache will go to a node in a different Availability Zone, meaning latency will be slightly higher. For more details, refer to Creating a Cache Cluster in the Amazon ElastiCache User Guide.
As mentioned at the outset, ElastiCache can be coupled with a wide variety of databases. Here is an example architecture that uses Amazon DynamoDB instead of Amazon RDS and MySQL:

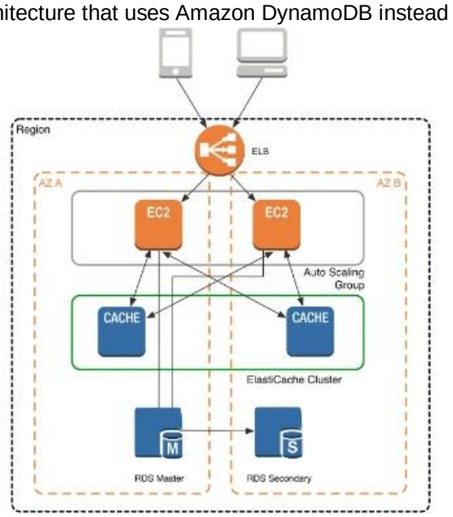

This combination of DynamoDB and ElastiCache is very popular with mobile and game companies, because DynamoDB allows for higher write throughput at lower cost than traditional relational databases. In addition, DynamoDB uses a key- value access pattern similar to ElastiCache, which also simplifies the programming model. Instead of using relational SQL for the primary database but then key-value patterns for the cache, both the primary database and cache can be programmed similarly. In this architecture pattern, DynamoDB remains the source of truth for data, but application reads are offloaded to ElastiCache for a speed boost.

QUESTION 14
An ERP application is deployed across multiple AZs in a single region. In the event of failure, the Recovery Time Objective (RTO) must be less than 3 hours, and the Recovery Point Objective (RPO) must be 15 minutes the customer realizes that data corruption occurred roughly 1.5 hours ago.
What DR strategy could be used to achieve this RTO and RPO in the event of this kind of failure?

A. Take hourly DB backups to S3, with transaction logs stored in S3 every 5 minutes.
B. Use synchronous database master-slave replication between two availability zones.
C. Take hourly DB backups to EC2 Instance store volumes with transaction logs stored In S3 every 5 minutes.

D. Take 15 minute DB backups stored In Glacier with transaction logs stored in S3 every 5 minutes.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 15
You are designing the network infrastructure for an application server in Amazon VPC. Users will access all application instances from the Internet, as well as from an on-premises network. The on-premises network is connected to your VPC over an AWS Direct Connect link.
How would you design routing to meet the above requirements?

A. Configure a single routing table with a default route via the Internet gateway. Propagate a default route via BGP on the AWS Direct Connect customer router. Associate the routing table with all VPC subnets.
B. Configure a single routing table with a default route via the Internet gateway. Propagate specific routes for the on-premises networks via BGP on the AWS Direct Connect customer router. Associate the routing table with all VPC subnets.
C. Configure a single routing table with two default routes: on to the Internet via an Internet gateway, the other to the on-premises network via the VPN gateway, Use this routing table across all subnets in the VPC.
D. Configure two routing tables: on that has a default router via the Internet gateway, and other that has a default route via the VPN gateway. Associate both routing tables with each VPC subnet.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 16
You control access to S3 buckets and objects with:

A. Identity and Access Management (IAM) Policies
B. Access Control Lists (ACLs).
C. Bucket Policies.
D. All of the above

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 17
The AWS IT infrastructure that AWS provides, complies with the following IT security standards, including:

A. SOC 1/SSAE 16/ISAE 3402 (formerly SAS 70 Type II), SOC 2 and SOC 3
B. FISMA, DIACAP, and FedRAMP
C. PCI DSS Level 1, ISO 27001, ITAR and FIPS 140-2
D. HIPAA, Cloud Security Alliance (CSA) and Motion Picture Association of America (MPAA)
E. All of the above

Correct Answer: ABC
Section: (none)
Explanation

Explanation/Reference:

QUESTION 18
Auto Scaling requests are signed with a_____signature calculated from the request and the user's private key.

A. SSL
B. AES-256
C. HMAC-SHA1
D. X.509

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 19
The following policy can be attached to an IAM group. It lets an IAM user in that group access a "home directory" in AWS S3 that matches their user name using the console.

```
{
"Version": "2012-10-17",
"Statement": [
{
"Action": ["s3:*"],
"Effect": "Allow",
"Resource": ["arn:aws:s3:::bucket-name"],
"Condition":{"StringLike":{"s3:prefix":["home/${aws:username}/*"]}}
},
{
"Action":["s3:*"],
"Effect":"Allow",
"Resource": ["arn:aws:s3:::bucket-name/home/${aws:username}/*"]
}
]
}
```

A. True
B. False

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 20
What does elasticity mean to AWS?

A. The ability to scale computing resources up easily, with minimal friction and down with latency
B. The ability to scale computing resources up and down easily, with minimal friction.
C. The ability to provision cloud computing resources in expectation of future demand.
D. The ability to recover from business continuity events with minimal friction.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 21
The following are AWS Storage services? Choose 2 Answers

A. AWS Relational Database Service (AWS RDS)

B. AWS ElastiCache
C. AWS Glacier
D. AWS ImporVExport

Correct Answer: BD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 22
How is AWS readily distinguished from other vendors in the traditional IT computing landscape?

A. Experienced. Scalable and elastic. Secure. Cost-effective. Reliable
B. Secure. Flexible. Cost-effective. Scalable and elastic. Global
C. Secure. Flexible. Cost-effective. Scalable and elastic. Experienced
D. Flexible. Cost-effective. Dynamic. Secure. Experienced.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 23
You have launched an EC2 instance with four (4) 500 GB EBS Provisioned IOPS volumes attached. The EC2 instance is EBS-Optimized and supports 500 Mbps throughput between EC2 and EBS. The four EBS volumes are configured as a single RAID 0 device, and each Provisioned IOPS volume is provisioned with 4,000 IOPS (4,000 16KB reads or writes), for a total of 16,000 random IOPS on the instance. The EC2 instance initially delivers the expected 16,000 IOPS random read and write performance. Sometime later, in order to increase the total random I/O performance of the instance, you add an additional two 500 GB EBS Provisioned IOPS volumes to the RAID. Each volume is provisioned to 4,000 IOPs like the original four, for a total of 24,000 IOPS on the EC2 instance. Monitoring shows that the EC2 instance CPU utilization increased from 50% to 70%, but the total random IOPS measured at the instance level does not increase at all.
What is the problem and a valid solution?

A. The EBS-Optimized throughput limits the total IOPS that can be utilized; use an EBSOptimized instance that provides larger throughput.
B. Small block sizes cause performance degradation, limiting the I/O throughput; configure the instance device driver and filesystem to use 64KB blocks to increase throughput.
C. The standard EBS Instance root volume limits the total IOPS rate; change the instance root volume to also be a 500GB 4,000 Provisioned IOPS volume.
D. Larger storage volumes support higher Provisioned IOPS rates; increase the provisioned volume storage of each of the 6 EBS volumes to 1TB.
E. RAID 0 only scales linearly to about 4 devices; use RAID 0 with 4 EBS Provisioned IOPS volumes, but increase each Provisioned IOPS EBS volume to 6,000 IOPS.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 24
Your company is storing millions of sensitive transactions across thousands of 100-GB files that must be encrypted in transit and at rest. Analysts concurrently depend on subsets of files, which can consume up to 5 TB of space, to generate simulations that can be used to steer business decisions. You are required to design an AWS solution that can cost effectively accommodate the long-term storage and in-flight subsets of data.

A. Use Amazon Simple Storage Service (S3) with server-side encryption, and run simulations on subsets in ephemeral drives on Amazon EC2.
B. Use Amazon S3 with server-side encryption, and run simulations on subsets in-memory on Amazon EC2.
C. Use HDFS on Amazon EMR, and run simulations on subsets in ephemeral drives on Amazon EC2.
D. Use HDFS on Amazon Elastic MapReduce (EMR), and run simulations on subsets in-memory on Amazon Elastic Compute Cloud (EC2).
E. Store the full data set in encrypted Amazon Elastic Block Store (EBS) volumes, and regularly capture snapshots that can be cloned to EC2 workstations.

Correct Answer: D
Section: (none)

Explanation

Explanation/Reference:

QUESTION 25
Your customer is willing to consolidate their log streams (access logs application logs security logs etc.) in one single system. Once consolidated, the customer wants to analyze these logs in real time based on heuristics. From time to time, the customer needs to validate heuristics, which requires going back to data samples extracted from the last 12 hours?
What is the best approach to meet your customer's requirements?

A. Send all the log events to Amazon SQS, setup an Auto Scaling group of EC2 servers to consume the logs and apply the heuristics.
B. Send all the log events to Amazon Kinesis, develop a client process to apply heuristics on the logs
C. Configure Amazon CloudTrail to receive custom logs, use EMR to apply heuristics the logs
D. Setup an Auto Scaling group of EC2 syslogd servers, store the logs on S3, use EMR to apply heuristics on the logs

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The throughput of an Amazon Kinesis stream is designed to scale without limits via increasing the number of shards within a stream. However, there are certain limits you should keep in mind while using Amazon Kinesis Streams:
By default, Records of a stream are accessible for up to 24 hours from the time they are added to the stream. You can raise this limit to up to 7 days by enabling extended data retention.
The maximum size of a data blob (the data payload before Base64-encoding) within one record is 1 megabyte (MB).
Each shard can support up to 1000 PUT records per second.
For more information about other API level limits, see Amazon Kinesis Streams Limits.

QUESTION 26
A newspaper organization has a on-premises application which allows the public to search its back catalogue and retrieve individual newspaper pages via a website written in Java They have scanned the old newspapers into JPEGs (approx 17TB) and used Optical Character Recognition (OCR) to populate a commercial search product. The hosting platform and software are now end of life and the organization wants to migrate Its archive to AWS and produce a cost efficient architecture and still be designed for availability and durability.
Which is the most appropriate?

A. Use S3 with reduced redundancy Io store and serve the scanned files, install the commercial search application on EC2 Instances and configure with auto-scaling and an Elastic Load Balancer.
B. Model the environment using CloudFormation use an EC2 instance running Apache webserver and an open source search application, stripe multiple standard EBS volumes together to store the JPEGs and search index.
C. Use S3 with standard redundancy to store and serve the scanned files, use CloudSearch for query processing, and use Elastic Beanstalk to host the website across multiple availability zones.
D. Use a single-AZ RDS MySQL instance Io store the search index 33d the JPEG images use an EC2 instance to serve the website and translate user queries into SQL.
E. Use a CloudFront download distribution to serve the JPEGs to the end users and Install the current commercial search product, along with a Java Container Tor the website on EC2 instances and use Route53 with DNS round-robin.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
There is no such thing as "Most appropriate" without knowing all your goals. I find your scenarios very fuzzy, since you
can obviously mix-n-match between them. I think you should decide by layers instead:
Load Balancer Layer: ELB or just DNS, or roll-your-own. (Using DNS+EIPs is slightly cheaper, but less reliable than ELB.)
Storage Layer for 17TB of Images: This is the perfect use case for S3. Off-load all the web requests directly to the
relevant JPEGs in S3. Your EC2 boxes just generate links to them.
If your app already serves it's own images (not links to images), you might start with EFS. But more than likely, you can
just setup a web server to re-write or re-direct all JPEG links to S3 pretty easily.
If you use S3, don't serve directly from the bucket - Serve via a CNAME in domain you control. That way, you can switch in CloudFront easily.
EBS will be way more expensive, and you'll need 2x the drives if you need 2 boxes. Yuck.
Consider a smaller storage format. For example, JPEG200 or WebP or other tools might make for smaller images. There is also the DejaVu format from a while back.
Cache Layer: Adding CloudFront in front of S3 will help people on the other side of the world -- well, possibly. Typical
archives follow a power law. The long tail of requests means that most JPEGs won't be requested enough to be in the
cache. So you are only speeding up the most popular objects. You can always wait, and switch in CF later after you know your costs better. (In some cases, it can actually lower costs.)
You can also put CloudFront in front of your app, since your archive search results should be fairly static. This will also
allow you to run with a smaller instance type, since CF will handle much of the load if you do it right.
Database Layer: A few options:

Use whatever your current server does for now, and replace with something else down the road. Don't under-estimate this approach, sometimes it's better to start now and optimize later.
Use RDS to run MySQL/Postgres
I'm not as familiar with ElasticSearch / Cloudsearch, but obviously Cloudsearch will be less maintenance+setup.
App Layer:
When creating the app layer from scratch, consider CloudFormation and/or OpsWorks. It's extra stuff to learn, but helps down the road.
Java+Tomcat is right up the alley of ElasticBeanstalk. (Basically EC2 + Autoscale + ELB).
Preventing Abuse: When you put something in a public S3 bucket, people will hot-link it from their web pages. If you
want to prevent that, your app on the EC2 box can generate signed links to S3 that expire in a few hours. Now everyone will be forced to go thru the app, and the app can apply rate limiting, etc.
Saving money: If you don't mind having downtime:
run everything in one AZ (both DBs and EC2s). You can always add servers and AZs down the road, as long as it's
architected to be stateless. In fact, you should use multiple regions if you want it to be really robust.
use Reduced Redundancy in S3 to save a few hundred bucks per month (Someone will have to "go fix it" every time it
breaks, including having an off-line copy to repair S3.)
Buy Reserved Instances on your EC2 boxes to make them cheaper. (Start with the RI market and buy a partially used one to get started.) It's just a coupon saying "if you run this type of box in this AZ, you will save on the per-hour costs." You can get 1/2 to 1/3 off easily.
Rewrite the application to use less memory and CPU - that way you can run on fewer/smaller boxes. (May or may not be worth the investment.)
If your app will be used very infrequently, you will save a lot of money by using Lambda. I'd be worried that it would be
quite slow if you tried to run a Java application on it though.
We're missing some information like load, latency expectations from search, indexing speed, size of the search index,
etc. But with what you've given us, I would go with S3 as the storage for the files (S3 rocks. It is really, really awesome). If you're stuck with the commercial search application, then on EC2 instances with autoscaling and an ELB. If you are allowed an alternative search engine, Elasticsearch is probably your best bet. I'd run it on EC2 instead of the AWS Elasticsearch service, as IMHO it's not ready yet. Don't autoscale Elasticsearch automatically though, it'll cause all sorts of issues. I have zero experience with CloudSearch so ic an't comment on that. Regardless of which option, I'd use CloudFormation for all of it.

QUESTION 27
Your company has recently extended its datacenter into a VPC on AWS to add burst computing capacity as needed Members of your Network Operations Center need to be able to go to the AWS Management Console and administer Amazon EC2 instances as necessary You don't want to create new IAM users for each NOC member and make those users sign in again to the AWS Management Console.
Which option below will meet the needs for your NOC members?

A. Use OAuth 2 0 to retrieve temporary AWS security credentials to enable your NOC members to sign in to the AWS Management Console.
B. Use web Identity Federation to retrieve AWS temporary security credentials to enable your NOC members to sign in to the AWS Management Console.
C. Use your on-premises SAML 2.0-compliant identity provider (IDP) to grant the NOC members federated access to the AWS Management Console via the AWS single sign-on (SSO) endpoint.
D. Use your on-premises SAML 2.0-compliam identity provider (IDP) to retrieve temporary security credentials to enable NOC members to sign in to the AWS Management Console.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 28
You are looking to migrate your Development (Dev) and Test environments to AWS. You have decided to use separate AWS accounts to host each environment. You plan to link each accounts bill to a Master AWS account using Consolidated Billing. To make sure you Keep within budget you would like to implement a way for administrators in the Master account to have access to stop, delete and/or terminate resources in both the Dev and Test accounts.
Identify which option will allow you to achieve this goal.

A. Create IAM users in the Master account with full Admin permissions. Create cross-account roles in the Dev and Test accounts that grant the Master account access to the resources in the account by inheriting permissions from the Master account.
B. Create IAM users and a cross-account role in the Master account that grants full Admin permissions to the Dev and Test accounts.
C. Create IAM users in the Master account Create cross-account roles in the Dev and Test accounts that have full Admin permissions and grant the Master account access.
D. Link the accounts using Consolidated Billing. This will give IAM users in the Master account access to resources in the Dev and Test accounts

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Bucket Owner Granting Cross-account Permission to objects It Does Not Own
In this example scenario, you own a bucket and you have enabled other AWS accounts to upload objects. That is, your
bucket can have objects that other AWS accounts own.
Now, suppose as a bucket owner, you need to grant cross-account permission on objects, regardless of who the owner

is, to a user in another account. For example, that user could be a billing application that needs to access object
metadata. There are two core issues:
The bucket owner has no permissions on those objects created by other AWS accounts. So for the bucket owner to grant permissions on objects it does not own, the object owner, the AWS account that created the objects, must first grant permission to the bucket owner. The bucket owner can then delegate those permissions.
Bucket owner account can delegate permissions to users in its own account but it cannot delegate permissions to other AWS accounts, because cross-account delegation is not supported.
In this scenario, the bucket owner can create an AWS Identity and Access Management (IAM) role with permission to
access objects, and grant another AWS account permission to assume the role temporarily enabling it to access objects in the bucket.
Background: Cross-Account Permissions and Using IAM Roles
IAM roles enable several scenarios to delegate access to your resources, and cross-account access is one of the key
scenarios. In this example, the bucket owner, Account A, uses an IAM role to temporarily delegate object access cross¬account to users in another AWS account, Account C. Each IAM role you create has two policies attached to it:
A trust policy identifying another AWS account that can assume the role.
An access policy defining what permissions—for example, s3:GetObject—are allowed when someone assumes the role. For a list of permissions you can specify in a policy, see Specifying Permissions in a Policy.
The AWS account identified in the trust policy then grants its user permission to assume the role. The user can then do the following to access objects:
Assume the role and, in response, get temporary security credentials.
Using the temporary security credentials, access the objects in the bucket.
For more information about IAM roles, go to Roles (Delegation and Federation) in IAM User Guide.
The following is a summary of the walkthrough steps:

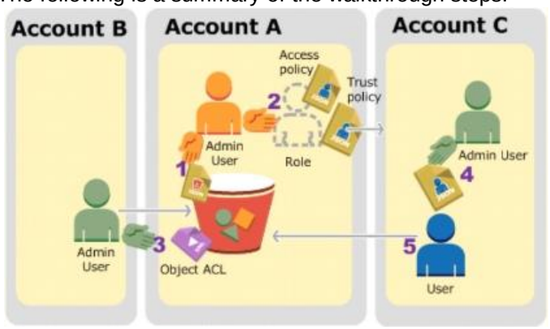

Account A administrator user attaches a bucket policy granting Account B conditional permission to upload objects.
Account A administrator creates an IAM role, establishing trust with Account C, so users in that account can access
Account A. The access policy attached to the role limits what user in Account C can do when the user accesses Account A.
Account B administrator uploads an object to the bucket owned by Account A, granting full-control permission to the
bucket owner.
Account C administrator creates a user and attaches a user policy that allows the user to assume the role.
User in Account C first assumes the role, which returns the user temporary security credentials. Using those temporary credentials, the user then accesses objects in the bucket.
For this example, you need three accounts. The following table shows how we refer to these accounts and the
administrator users in these accounts. Per IAM guidelines (see About Using an Administrator User to Create Resources
and Grant Permissions) we do not use the account root credentials in this walkthrough. Instead, you create an
administrator user in each account and use those credentials in creating resources and granting them permissions

AWS Account ID	Account Referred To As	Administrator User in the Account
1111-1111-1111	Account A	AccountAadmin
2222-2222-2222	Account B	AccountBadmin
3333-3333-3333	Account C	AccountCadmin

QUESTION 29
You're running an application on-premises due to its dependency on non-x86 hardware and want to use AWS for data backup. Your backup application is only able to write to POSIX-compatible block-based storage. You have 140TB of data and would like to mount it as a single folder on your file server Users must be able to access portions of this data while the backups are taking place.
What backup solution would be most appropriate for this use case?

A. Use Storage Gateway and configure it to use Gateway Cached volumes.
B. Configure your backup software to use S3 as the target for your data backups.
C. Configure your backup software to use Glacier as the target for your data backups.
D. Use Storage Gateway and configure it to use Gateway Stored volumes.

Correct Answer: A

Section: (none)
Explanation

Explanation/Reference:
Explanation:
Gateway-Cached Volume Architecture
Gateway-cached volumes let you use Amazon Simple Storage Service (Amazon S3) as your primary data storage while
retaining frequently accessed data locally in your storage gateway. Gateway-cached volumes minimize the need to scale your on-premises storage infrastructure, while still providing your applications with low-latency access to their frequently accessed data. You can create storage volumes up to 32 TiB in size and attach to them as iSCSI devices from your on-premises application servers. Your gateway stores data that you write to these volumes in Amazon S3 and retains recently read data in your on-premises storage gateway's cache and upload buffer storage.
Gateway-cached volumes can range from 1 GiB to 32 TiB in size and must be rounded to the nearest GiB. Each gateway configured for gateway-cached volumes can support up to 32 volumes for a total maximum storage volume of 1,024 TiB (1 PiB).
In the gateway-cached volume solution, AWS Storage Gateway stores all your on-premises application data in a storage volume in Amazon S3.
The following diagram provides an overview of the AWS Storage Gateway-cached volume deployment.

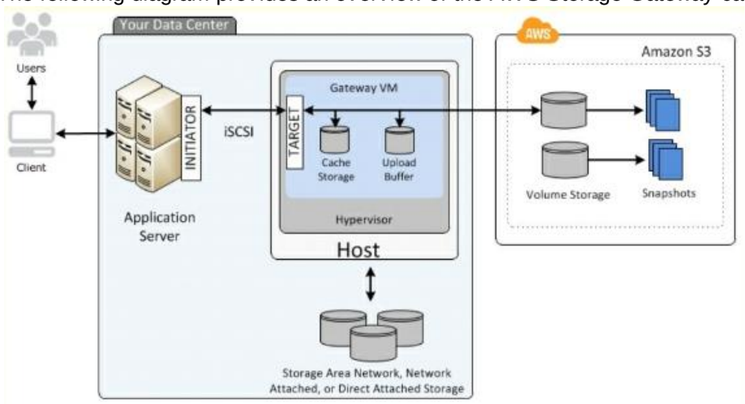

After you've installed the AWS Storage Gateway software appliance—the virtual machine (VM)—on a host in your data center and activated it, you can use the AWS Management Console to provision storage volumes backed by Amazon S3. You can also provision storage volumes programmatically using the AWS Storage Gateway API or the AWS SDK libraries. You then mount these storage volumes to your on-premises application servers as iSCSI devices.
You also allocate disks on-premises for the VM. These on-premises disks serve the following purposes:
Disks for use by the gateway as cache storage - As your applications write data to the storage volumes in AWS, the gateway initially stores the data on the on-premises disks referred to as cache storage before uploading the data to Amazon S3. The cache storage acts as the on-premises durable store for data that is waiting to upload to Amazon S3 from the upload buffer.
The cache storage also lets the gateway store your application's recently accessed data on-premises for low-latency access. If your application requests data, the gateway first checks the cache storage for the data before checking Amazon S3.
You can use the following guidelines to determine the amount of disk space to allocate for cache storage. Generally, you should allocate at least 20 percent of your existing file store size as cache storage. Cache storage should also be larger than the upload buffer. This latter guideline helps ensure cache storage is large enough to persistently hold all data in the upload buffer that has not yet been uploaded to Amazon S3.
Disks for use by the gateway as the upload buffer - To prepare for upload to Amazon S3, your gateway also stores incoming data in a staging area, referred to as an upload buffer. Your gateway uploads this buffer data over an encrypted Secure Sockets Layer (SSL) connection to AWS, where it is stored encrypted in Amazon S3.
You can take incremental backups, called snapshots, of your storage volumes in Amazon S3. These point-in-time snapshots are also stored in Amazon S3 as Amazon EBS snapshots. When you take a new snapshot, only the data that has changed since your last snapshot is stored. You can initiate snapshots on a scheduled or one-time basis. When you delete a snapshot, only the data not needed for any other snapshots is removed.
You can restore an Amazon EBS snapshot to a gateway storage volume if you need to recover a backup of your data. Alternatively, for snapshots up to 16 TiB in size, you can use the snapshot as a starting point for a new Amazon EBS volume. You can then attach this new Amazon EBS volume to an Amazon EC2 instance.
All gateway-cached volume data and snapshot data is stored in Amazon S3 encrypted at rest using server-side
encryption (SSE). However, you cannot access this data with the Amazon S3 API or other tools such as the Amazon S3 console.

QUESTION 30
To serve Web traffic for a popular product your chief financial officer and IT director have purchased 10 ml large heavy utilization Reserved Instances (RIs) evenly spread across two availability zones: Route 53 is used to deliver the traffic to an Elastic Load Balancer (ELB). After several months, the product grows even more popular and you need additional capacity. As a result, your company purchases two C3.2xlarge medium utilization Ris. You register the two c3 2xlarge instances with your ELB and quickly find that the ml large instances are at 100% of capacity and the c3 2xlarge instances have significant capacity that's unused.
Which option is the most cost effective and uses EC2 capacity most effectively?

A. Configure Autoscaling group and Launch Configuration with ELB to add up to 10 more on-demand ml .large instances when triggered by Cloudwatch. Shut off c3.2xlarge instances.
B. Configure ELB with two c3.2xlarge instances and use on-demand Autoscaling group for up to two additional c3.2xlarge instances. Shut off ml .large instances.
C. Route traffic to EC2 ml .large and c3.2xlarge instances directly using Route 53 latency based routing and health checks. Shut off ELB.
D. Use a separate ELB for each instance type and distribute load to ELBs with Route 53 weighted round robin.

Correct Answer: B
Section: (none)

Explanation

Explanation/Reference:

QUESTION 31
You have deployed a web application targeting a global audience across multiple AWS Regions under the domain name.example.com. You decide to use Route53 Latency-Based Routing to serve web requests to users from the region closest to the user. To provide business continuity in the event of server downtime you configure weighted record sets associated with two web servers in separate Availability Zones per region. Dunning a DR test you notice that when you disable all web servers in one of the regions Route53 does not automatically direct all users to the other region.
What could be happening? (Choose 2 answers)

A. Latency resource record sets cannot be used in combination with weighted resource record sets.
B. You did not setup an HTTP health check to one or more of the weighted resource record sets associated with me disabled web servers.
C. The value of the weight associated with the latency alias resource record set in the region with the disabled servers is higher than the weight for the other region.
D. One of the two working web servers in the other region did not pass its HTTP health check.
E. You did not set "Evaluate Target Health" to "Yes" on the latency alias resource record set associated with example com in the region where you disabled the servers.

Correct Answer: BE
Section: (none)
Explanation

Explanation/Reference:
Explanation:
How Health Checks Work in Complex Amazon Route 53 Configurations
Checking the health of resources in complex configurations works much the same way as in simple configurations. However, in complex configurations, you use a combination of alias resource record sets (including weighted alias, latency alias, and failover alias) and nonalias resource record sets to build a decision tree that gives you greater control over how Amazon Route 53 responds to requests. For more information, see How Health Checks Work in Simple Amazon Route 53 Configurations.
For example, you might use latency alias resource record sets to select a region close to a user and use weighted resource record sets for two or more resources within each region to protect against the failure of a single endpoint or an Availability Zone. The following diagram shows this configuration.

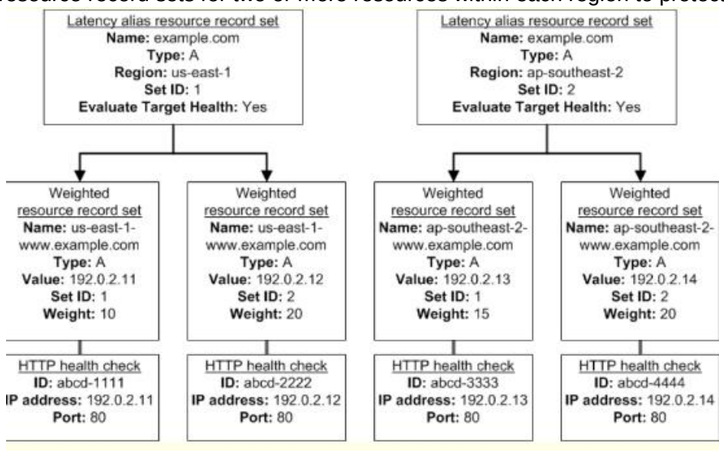

Here's how Amazon EC2 and Amazon Route 53 are configured:
You have Amazon EC2 instances in two regions, us-east-1 and ap-southeast-2. You want Amazon Route 53 to respond to queries by using the resource record sets in the region that provides the lowest latency for your customers, so you create a latency alias resource record set for each region. (You create the latency alias resource record sets after you create resource record sets for the individual Amazon EC2 instances.)
Within each region, you have two Amazon EC2 instances. You create a weighted resource record set for each instance.
The name and the type are the same for both of the weighted resource record sets in each region.
When you have multiple resources in a region, you can create weighted or failover resource record sets for your resources. You can also create even more complex configurations by creating weighted alias or failover alias resource record sets that, in turn, refer to multiple resources.
Each weighted resource record set has an associated health check. The IP address for each health check matches the IP address for the corresponding resource record set. This isn't required, but it's the most common configuration.
For both latency alias resource record sets, you set the value of Evaluate Target Health to Yes.
You use the Evaluate Target Health setting for each latency alias resource record set to make Amazon Route 53 evaluate the health of the alias targets—the weighted resource record sets—and respond accordingly.

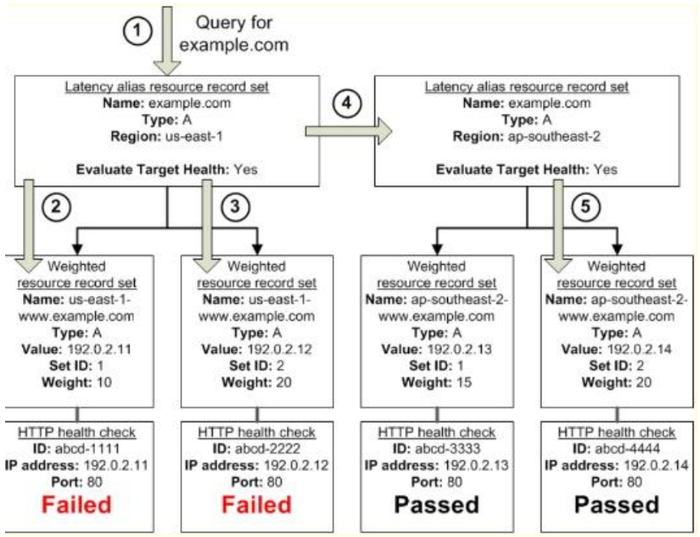

The preceding diagram illustrates the following sequence of events:
Amazon Route 53 receives a query for example.com. Based on the latency for the user making the request, Amazon Route 53 selects the latency alias resource record set for the us-east-1 region.
Amazon Route 53 selects a weighted resource record set based on weight. Evaluate Target Health is Yes for the latency alias resource record set, so Amazon Route 53 checks the health of the selected weighted resource record set.
The health check failed, so Amazon Route 53 chooses another weighted resource record set based on weight and checks its health. That resource record set also is unhealthy.
Amazon Route 53 backs out of that branch of the tree, looks for the latency alias resource record set with the next-best latency, and chooses the resource record set for ap-southeast-2.
Amazon Route 53 again selects a resource record set based on weight, and then checks the health of the selected resource record set. The health check passed, so Amazon Route 53 returns the applicable value in response to the query.
What Happens When You Associate a Health Check with an Alias Resource Record Set?
You can associate a health check with an alias resource record set instead of or in addition to setting the value of Evaluate Target Health to Yes. However, it's generally more useful if Amazon Route 53 responds to queries based on the health of the underlying resources—the HTTP servers, database servers, and other resources that your alias resource record sets refer to. For example, suppose the following configuration:
You assign a health check to a latency alias resource record set for which the alias target is a group of weighted resource record sets.
You set the value of Evaluate Target Health to Yes for the latency alias resource record set.
In this configuration, both of the following must be true before Amazon Route 53 will return the applicable value for a weighted resource record set:
The health check associated with the latency alias resource record set must pass.
At least one weighted resource record set must be considered healthy, either because it's associated with a health check that passes or because it's not associated with a health check. In the latter case, Amazon Route 53 always considers the weighted resource record set healthy.

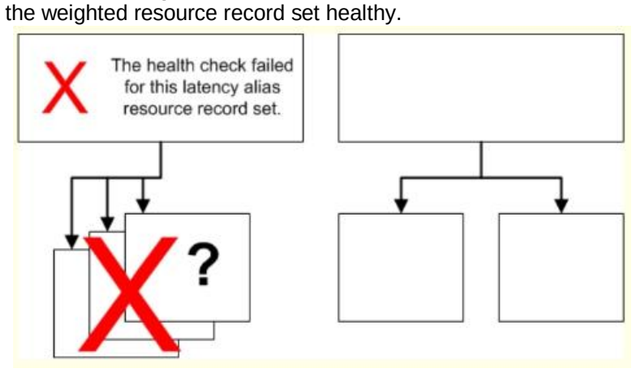

If the health check for the latency alias resource record set fails, Amazon Route 53 stops responding to queries using any of the weighted resource record sets in the alias target, even if they're all healthy. Amazon Route 53 doesn't know the status of the weighted resource record sets because it never looks past the failed health check on the alias resource
record set.
What Happens When You Omit Health Checks?
In a complex configuration, it's important to associate health checks with all of the non-alias resource record sets. Let's
return to the preceding example, but assume that a health check is missing on one of the weighted resource record sets in the us-east-1 region:

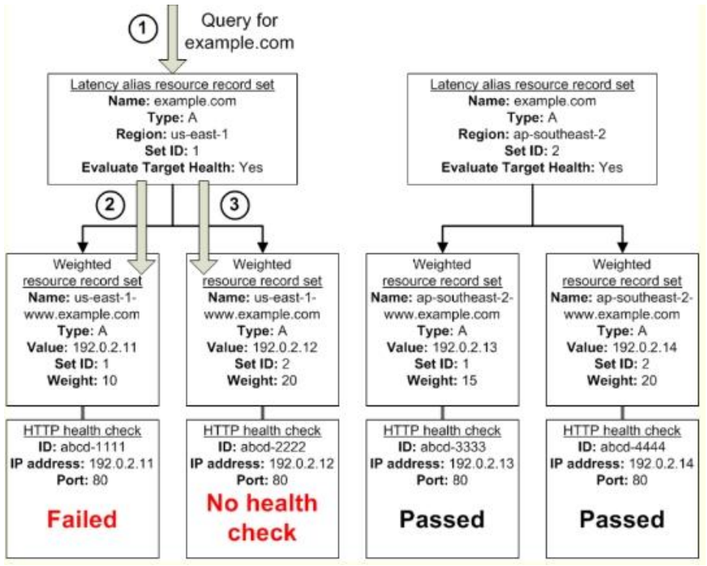

Here's what happens when you omit a health check on a non-alias resource record set in this configuration:
Amazon Route 53 receives a query for example.com. Based on the latency for the user making the request, Amazon Route 53 selects the latency alias resource record set for the us-east-1 region.
Amazon Route 53 looks up the alias target for the latency alias resource record set, and checks the status of the corresponding health checks. The health check for one weighted resource record set failed, so that resource record set is omitted from consideration.
The other weighted resource record set in the alias target for the us-east-1 region has no health check. The corresponding resource might or might not be healthy, but without a health check, Amazon Route 53 has no way to know. Amazon Route 53 assumes that the resource is healthy and returns the applicable value in response to the query. What Happens When You Set Evaluate Target Health to No?
In general, you also want to set Evaluate Target Health to Yes for all of the alias resource record sets. In the following example, all of the weighted resource record sets have associated health checks, but Evaluate Target Health is set to No for the latency alias resource record set for the us-east-1 region:

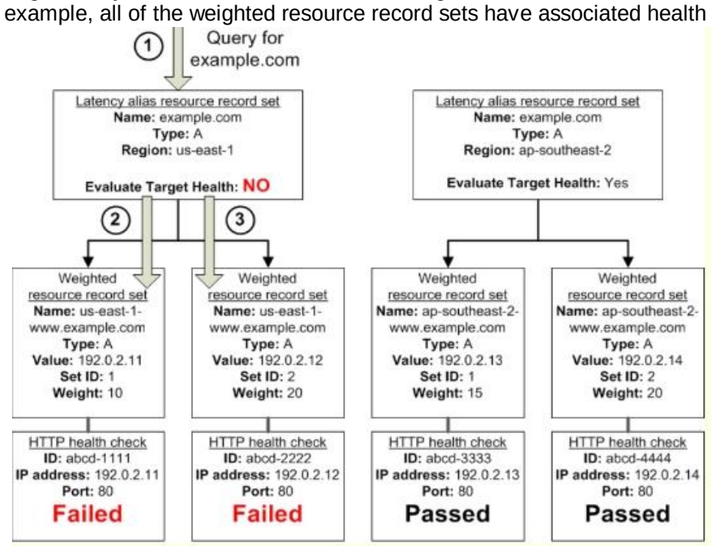

Here's what happens when you set Evaluate Target Health to No for an alias resource record set in this configuration:
Amazon Route 53 receives a query for example.com. Based on the latency for the user making the request, Amazon Route 53 selects the latency alias resource record set for the us-east-1 region.
Amazon Route 53 determines what the alias target is for the latency alias resource record set, and checks the corresponding health checks. They're both failing.
Because the value of Evaluate Target Health is No for the latency alias resource record set for the us-east-1 region, Amazon Route 53 must choose one resource record set in this branch instead of backing out of the branch and looking for a healthy resource record set in the ap-southeast-2 region.

QUESTION 32
Your startup wants to implement an order fulfillment process for selling a personalized gadget that needs an average of 3-4 days to produce with some orders taking up to 6 months you expect 10 orders per day on your first day. 1000 orders per day after 6 months and 10,000 orders after 12 months.

Orders coming in are checked for consistency men dispatched to your manufacturing plant for production quality control packaging shipment and payment processing If the product does not meet the quality standards at any stage of the process employees may force the process to repeat a step Customers are notified via email about order status and any critical issues with their orders such as payment failure.
Your case architecture includes AWS Elastic Beanstalk for your website with an RDS MySQL instance for customer data and orders.
How can you implement the order fulfillment process while making sure that the emails are delivered reliably?

A. Add a business process management application to your Elastic Beanstalk app servers and re-use the ROS database for tracking order status use one of the Elastic Beanstalk instances to send emails to customers.
B. Use SWF with an Auto Scaling group of activity workers and a decider instance in another Auto Scaling group with min/max=1 Use the decider instance to send emails to customers.
C. Use SWF with an Auto Scaling group of activity workers and a decider instance in another Auto Scaling group with min/max=1 use SES to send emails to customers.
D. Use an SQS queue to manage all process tasks Use an Auto Scaling group of EC2 Instances that poll the tasks and execute them. Use SES to send emails to customers.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 33
A read only news reporting site with a combined web and application tier and a database tier that receives large and unpredictable traffic demands must be able to respond to these traffic fluctuations automatically.
What AWS services should be used meet these requirements?

A. Stateless instances for the web and application tier synchronized using ElastiCache Memcached in an autoscaimg group monitored with CloudWatch and RDS with read replicas.
B. Stateful instances for the web and application tier in an autoscaling group monitored with CloudWatch and RDS with read replicas.
C. Stateful instances for the web and application tier in an autoscaling group monitored with Cloudwatch. And multi-AZ RDS.
D. Stateless instances for the web and application tier synchronized using ElastiCache Memcached in an autoscaling group monitored with Cloudwatch and multi-AZ RDS.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 34
You are designing a photo-sharing mobile app. The application will store all pictures in a single Amazon S3 bucket.
Users will upload pictures from their mobile device directly to Amazon S3 and will be able to view and download their own pictures directly from Amazon S3. You want to configure security to handle potentially millions of users in the most secure manner possible.
What should your server-side application do when a new user registers on the photo-sharing mobile application?

A. Create an IAM user. Update the bucket policy with appropriate permissions for the IAM user. Generate an access key and secret key for the IAM user, store them in the mobile app and use these credentials to access Amazon S3.
B. Create an IAM user. Assign appropriate permissions to the IAM user. Generate an access key and secret key for the IAM user, store them in the mobile app and use these credentials to access Amazon S3.
C. Create a set of long-term credentials using AWS Security Token Service with appropriate permissions. Store these credentials in the mobile app and use them to access Amazon S3.
D. Record the user's information in Amazon RDS and create a role in IAM with appropriate permissions. When the user uses their mobile app, create temporary credentials using the AWS Security Token Service "AssumeRole" function. Store these credentials in the mobile app's memory and use them to access Amazon S3. Generate new credentials the next time the user runs the mobile app.
E. Record the user's information in Amazon DynamoDB. When the user uses their mobile app, create temporary credentials using AWS Security Token Service with appropriate permissions. Store these credentials in the mobile app's memory and use them to access Amazon S3. Generate new credentials the next time the user runs the mobile app.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 35
You are tasked with moving a legacy application from a virtual machine running Inside your datacenter to an Amazon VPC Unfortunately this app requires access to a number of on-premises services and no one who configured the app still works for your company. Even worse there's no documentation for it. What will allow the application running inside the VPC to reach back and access its internal dependencies without being reconfigured? (Choose 3 answers)

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

B. An Internet Gateway to allow a VPN connection.
C. An Elastic IP address on the VPC instance
D. An IP address space that does not conflict with the one on-premises
E. Entries in Amazon Route 53 that allow the Instance to resolve its dependencies' IP addresses
F. A VM Import of the current virtual machine

Correct Answer: ADF
Section: (none)
Explanation

Explanation/Reference:
Explanation:
AWS Direct Connect
AWS Direct Connect makes it easy to establish a dedicated network connection from your premises to AWS. Using AWS Direct Connect, you can establish private connectivity between AWS and your datacenter, office, or collocation environment, which in many cases can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than Internet-based connections.
AWS Direct Connect lets you establish a dedicated network connection between your network and one of the AWS
Direct Connect locations. Using industry standard 802.1q VLANs, this dedicated connection can be partitioned into
multiple virtual interfaces. This allows you to use the same connection to access public resources such as objects stored in Amazon S3 using public IP address space, and private resources such as Amazon EC2 instances running within an Amazon Virtual Private Cloud (VPC) using private IP space, while maintaining network separation between the public and private environments. Virtual interfaces can be reconfigured at any time to meet your changing needs.
What is AWS Direct Connect?
AWS Direct Connect links your internal network to an AWS Direct Connect location over a standard 1 gigabit or 10 gigabit Ethernet fiber-optic cable. One end of the cable is connected to your router, the other to an AWS Direct Connect router. With this connection in place, you can create virtual interfaces directly to the AWS cloud (for example, to Amazon Elastic Compute Cloud (Amazon EC2) and Amazon Simple Storage Service (Amazon S3) and to Amazon Virtual Private Cloud (Amazon VPC), bypassing Internet service providers in your network path. An AWS Direct Connect location provides access to Amazon Web Services in the region it is associated with, as well as access to other US regions. For example, you can provision a single connection to any AWS Direct Connect location in the US and use it to access public AWS services in all US Regions and AWS GovCloud (US).
The following diagram shows how AWS Direct Connect interfaces with your network.

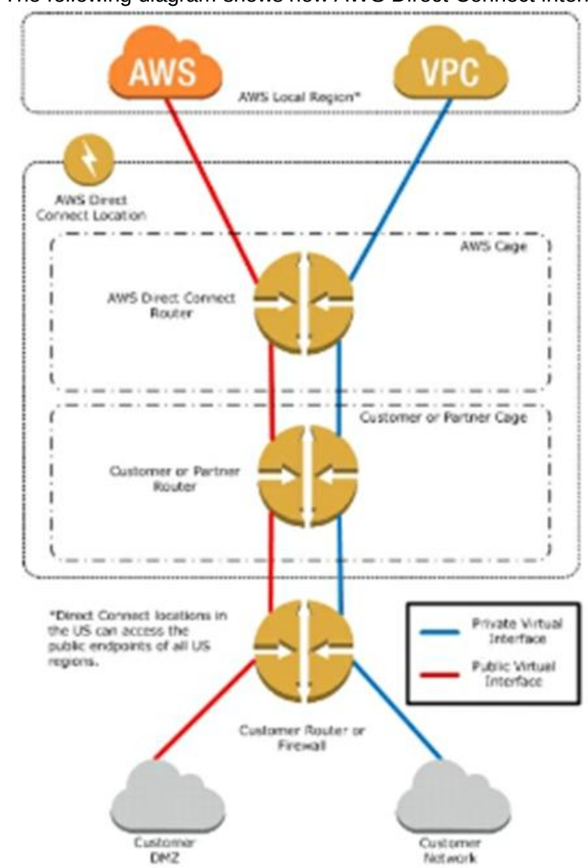

Requirements
To use AWS Direct Connect, your network must meet one of the following conditions:

Your network is collocated with an existing AWS Direct Connect location. For more information on available AWS Direct Connect locations, go to http://aws.amazon.com/directconnect/.
You are working with an AWS Direct Connect partner who is a member of the AWS Partner Network (APN). For a list of AWS Direct Connect partners who can help you connect, go to http://aws.amazon.com/directconnect.
You are working with an independent service provider to connect to AWS Direct Connect.
In addition, your network must meet the following conditions:
Connections to AWS Direct Connect require single mode fiber, 1000BASE-LX (1310nm) for 1 gigabit Ethernet, or
10GBASE-LR (1310nm) for 10 gigabit Ethernet. Auto Negotiation for the port must be disabled. You must support 802.1Q VLANs across these connections.
Your network must support Border Gateway Protocol (BGP) and BGP MD5 authentication. Optionally, you may configure Bidirectional Forwarding Detection (BFD).
To connect to Amazon Virtual Private Cloud (Amazon VPC), you must first do the following:
Provide a private Autonomous System Number (ASN). Amazon allocates a private IP address in the 169.x.x.x range to
you.
Create a virtual private gateway and attach it to your VPC. For more information about creating a virtual private
gateway, see Adding a Hardware Virtual Private Gateway to Your VPC in the Amazon VPC User Guide.
To connect to public AWS products such as Amazon EC2 and Amazon S3, you need to provide the following:
A public ASN that you own (preferred) or a private ASN.
Public IP addresses (/31) (that is, one for each end of the BGP session) for each BGP session. If you do not have public IP addresses to assign to this connection, log on to AWS and then open a ticket with AWS Support.
The public routes that you will advertise over BGP.

QUESTION 36
You have a periodic Image analysis application that gets some files In Input analyzes them and tor each file writes some data in output to a ten file the number of files in input per day is high and concentrated in a few hours of the day.
Currently you have a server on EC2 with a large EBS volume that hosts the input data and the results it takes almost 20 hours per day to complete the process.
What services could be used to reduce the elaboration time and improve the availability of the solution?

A. S3 to store I/O files. SQS to distribute elaboration commands to a group of hosts working in parallel. Auto scaling to dynamically size the group of hosts depending on the length of the SQS queue
B. EBS with Provisioned IOPS (PIOPS) to store I/O files. SNS to distribute elaboration commands to a group of hosts working in parallel Auto Scaling to dynamically size the group of hosts depending on the number of SNS notifications
C. S3 to store I/O files, SNS to distribute evaporation commands to a group of hosts working in parallel. Auto scaling to dynamically size the group of hosts depending on the number of SNS notifications
D. EBS with Provisioned IOPS (PIOPS) to store I/O files SQS to distribute elaboration commands to a group of hosts working in parallel Auto Scaling to dynamically size the group ot hosts depending on the length of the SQS queue.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Amazon EBS 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 a specific Availability Zone, where they are automatically replicated to protect you from the failure of a single component.
Amazon EBS provides three volume types: General Purpose (SSD), Provisioned IOPS (SSD), and Magnetic. The three volume types differ in performance characteristics and cost, so you can choose the right storage performance and price for the needs of your applications. All EBS volume types offer the same durable snapshot capabilities and are designed for 99.999% availability.

QUESTION 37
You have been asked to design the storage layer for an application. The application requires disk performance of at least 100,000 IOPS. In addition, the storage layer must be able to survive the loss of an individual disk, EC2 instance, or Availability Zone without any data loss. The volume you provide must have a capacity of at least 3 TB.
Which of the following designs will meet these objectives?

A. Instantiate a c3.8xlarge instance in us-east-1. Provision 4x1 TB EBS volumes, attach them to the instance, and configure them as a single RAID 5 volume. Ensure that EBS snapshots are performed every 15 minutes.
B. Instantiate a c3.8xlarge instance in us-east-1. Provision 3xlTB EBS volumes, attach them to the Instance, and configure them as a single RAID 0 volume. Ensure that EBS snapshots are performed every 15 minutes.
C. Instantiate an i2.8xlarge instance in us-east-la. Create a RAID 0 volume using the four 800GB SSD ephemeral disks provided with the instance. Provision 3x1TB EBS volumes, attach them to the instance, and configure them as a second RAID 0 volume. Configure synchronous, block-level replication from the ephemeral-backed volume to the EBS-backed volume.
D. Instantiate a c3.8xlarge instance in us-east-1. Provision an AWS Storage Gateway and configure it for 3 TB of storage and 100,000 IOPS. Attach the volume to the instance.
E. Instantiate an i2.8xlarge instance in us-east-la. Create a RAID 0 volume using the four 800GB SSD ephemeral disks provided with the instance. Configure synchronous, blocklevel replication to an identically configured instance in us-east-1 b.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 38
You are the new IT architect in a company that operates a mobile sleep tracking application.
When activated at night, the mobile app is sending collected data points of 1 kilobyte every 5 minutes to your backend.

The backend takes care of authenticating the user and writing the data points into an Amazon DynamoDB table.
Every morning, you scan the table to extract and aggregate last night's data on a per user basis, and store the results in Amazon S3. Users are notified via Amazon SNS mobile push notifications that new data is available, which is parsed and visualized by the mobile app.
Currently you have around 100k users who are mostly based out of North America.
You have been tasked to optimize the architecture of the backend system to lower cost.
What would you recommend? (Choose 2)

A. Have the mobile app access Amazon DynamoDB directly Instead of JSON files stored on Amazon S3.
B. Write data directly into an Amazon Redshift cluster replacing both Amazon DynamoDB and Amazon S3.
C. Introduce an Amazon SQS queue to buffer writes to the Amazon DynamoDB table and reduce provisioned write throughput.
D. Introduce Amazon Elasticache to cache reads from the Amazon DynamoDB table and reduce provisioned read throughput.
E. Create a new Amazon DynamoDB table each day and drop the one for the previous day after its data is on Amazon S3.

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 39
A large real-estate brokerage is exploring the option of adding a cost-effective location based alert to their existing mobile application. The application backend infrastructure currently runs on AWS Users who opt in to this service will receive alerts on their mobile device regarding real-estate otters in proximity to their location. For the alerts to be relevant delivery time needs to be in the low minute count the existing mobile app has 5 million users across the US.
Which one of the following architectural suggestions would you make to the customer?

A. The mobile application will submit its location to a web service endpoint utilizing Elastic Load Balancing and EC2 instances: DynamoDB will be used to store and retrieve relevant offers EC2 instances will communicate with mobile earners/device providers to push alerts back to mobile application.
B. Use AWS DirectConnect or VPN to establish connectivity with mobile carriers EC2 instances will receive the mobile applications ' location through carrier connection: RDS will be used to store and relevant offers EC2 instances will communicate with mobile carriers to push alerts back to the mobile application
C. The mobile application will send device location using SQS. EC2 instances will retrieve the relevant others from DynamoDB AWS Mobile Push will be used to send offers to the mobile application
D. The mobile application will send device location using AWS Mobile Push EC2 instances will retrieve the relevant offers from DynamoDB EC2 instances will communicate with mobile carriers/device providers to push alerts back to the mobile application.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 40
You currently operate a web application. In the AWS US-East region The application runs on an auto-scaled layer of EC2 instances and an RDS Multi-AZ database Your IT security compliance officer has tasked you to develop a reliable and durable logging solution to track changes made to your EC2.IAM And RDS resources.
The solution must ensure the integrity and confidentiality of your log dat

A. Which of these solutions would you recommend?
B. Create a new CloudTrail trail with one new S3 bucket to store the logs and with the global services option selected Use IAM roles S3 bucket policies and Multi Factor Authentication (MFA) Delete on the S3 bucket that stores your logs.
C. Create a new CloudTrail with one new S3 bucket to store the logs Configure SNS to send log file delivery notifications to your management system Use IAM roles and S3 bucket policies on the S3 bucket mat stores your logs.
D. Create a new CloudTrail trail with an existing S3 bucket to store the logs and with the global services option selected Use S3 ACLs and Multi Factor Authentication (MFA) Delete on the S3 bucket that stores your logs.
E. Create three new CloudTrail trails with three new S3 buckets to store the logs one for the AWS Management console, one for AWS SDKs and one for command line tools Use IAM roles and S3 bucket policies on the S3 buckets that store your logs.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 41

Your department creates regular analytics reports from your company's log files All log data is collected in Amazon S3 and processed by daily Amazon Elastic MapReduce (EMR) jobs that generate daily PDF reports and aggregated tables in CSV format for an Amazon Redshift data warehouse.
Your CFO requests that you optimize the cost structure for this system.
Which of the following alternatives will lower costs without compromising average performance of the system or data integrity for the raw data?

A. Use reduced redundancy storage (RRS) for all data In S3. Use a combination of Spot Instances and Reserved Instances for Amazon EMR jobs. Use Reserved Instances for Amazon Redshift.
B. Use reduced redundancy storage (RRS) for PDF and csv data in S3. Add Spot Instances to EMR jobs. Use Spot Instances for Amazon Redshift.
C. Use reduced redundancy storage (RRS) for PDF and .csv data In Amazon S3. Add Spot Instances to Amazon EMR jobs. Use Reserved Instances for Amazon Redshift.
D. Use reduced redundancy storage (RRS) for all data in Amazon S3. Add Spot Instances to Amazon EMR jobs. Use Reserved Instances for Amazon Redshift.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Using Reduced Redundancy Storage Amazon S3 stores objects according to their storage class. It assigns the storage class to an object when it is written to Amazon S3. You can assign objects a specific storage class (standard or reduced redundancy) only when you write the objects to an Amazon S3 bucket or when you copy objects that are already stored in Amazon S3. Standard is the default storage class. For information about storage classes, see Object Key and Metadata.
In order to reduce storage costs, you can use reduced redundancy storage for noncritical, reproducible data at lower levels of redundancy than Amazon S3 provides with standard storage. The lower level of redundancy results in less durability and availability, but in many cases, the lower costs can make reduced redundancy storage an acceptable storage solution. For example, it can be a cost-effective solution for sharing media content that is durably stored elsewhere. It can also make sense if you are storing thumbnails and other resized images that can be easily reproduced from an original image.
Reduced redundancy storage is designed to provide 99.99% durability of objects over a given year. This durability level corresponds to an average annual expected loss of 0.01% of objects. For example, if you store 10,000 objects using the RRS option, you can, on average, expect to incur an annual loss of a single object per year (0.01% of 10,000 objects).
Note:
This annual loss represents an expected average and does not guarantee the loss of less than 0.01% of objects in a given year.
Reduced redundancy storage stores objects on multiple devices across multiple facilities, providing 400 times the durability of a typical disk drive, but it does not replicate objects as many times as Amazon S3 standard storage. In addition, reduced redundancy storage is designed to sustain the loss of data in a single facility.
If an object in reduced redundancy storage has been lost, Amazon S3 will return a 405 error on requests made to that object. Amazon S3 also offers notifications for reduced redundancy storage object loss: you can configure your bucket so that when Amazon S3 detects the loss of an RRS object, a notification will be sent through Amazon Simple Notification Service (Amazon SNS). You can then replace the lost object. To enable notifications, you can use the Amazon S3 console to set the Notifications property of your bucket.

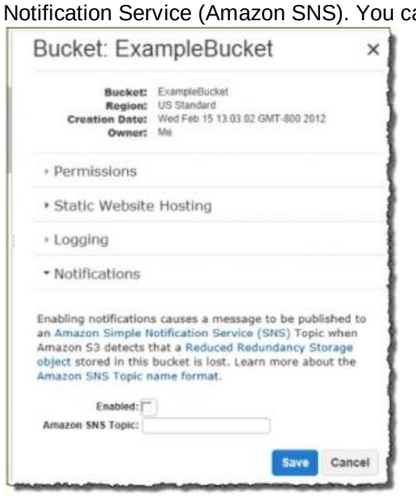

QUESTION 42
You require the ability to analyze a large amount of data, which is stored on Amazon S3 using Amazon Elastic Map Reduce. You are using the cc2 8x large Instance type, whose CPUs are mostly idle during processing. Which of the below would be the most cost efficient way to reduce the runtime of the job?

A. Create more smaller flies on Amazon S3.

B. Add additional cc2 8x large instances by introducing a task group.
C. Use smaller instances that have higher aggregate I/O performance.
D. Create fewer, larger files on Amazon S3.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 43
An AWS customer is deploying an application mat is composed of an AutoScaling group of EC2 Instances
The customers security policy requires that every outbound connection from these instances to any other service within the customers Virtual Private Cloud must be authenticated using a unique x 509 certificate that contains the specific instance-id.
In addition an x 509 certificates must Designed by the customer's Key management service in order to be trusted for authentication.
Which of the following configurations will support these requirements?

A. Configure an IAM Role that grants access to an Amazon S3 object containing a signed certificate and configure me Auto Scaling group to launch instances with this role Have the instances bootstrap get the certificate from Amazon S3 upon first boot.
B. Embed a certificate into the Amazon Machine Image that is used by the Auto Scaling group Have the launched instances generate a certificate signature request with the instance's assigned instance-id to the Key management service for signature.
C. Configure the Auto Scaling group to send an SNS notification of the launch of a new instance to the trusted key management service. Have the Key management service generate a signed certificate and send it directly to the newly launched instance.
D. Configure the launched instances to generate a new certificate upon first boot Have the Key management service poll the Auto Scaling group for associated instances and send new instances a certificate signature (hat contains the specific instance-id.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 44
Your company runs a customer facing event registration site This site is built with a 3-tier architecture with web and application tier servers and a MySQL database The application requires 5 web tier servers and 6 application tier servers for normal operation, but can run on a minimum of 65% server capacity and a single MySQL database.
When deploying this application in a region with three availability zones (AZs) which architecture provides high availability?

A. A web tier deployed across 2 AZs with 3 EC2 (Elastic Compute Cloud) instances in each AZ inside an Auto Scaling Group behind an ELB (elastic load balancer), and an application tier deployed across 2 AZs with 3 EC2 instances in each AZ inside an Auto Scaling Group behind an ELB and one RDS (Relational Database Service) instance deployed with read replicas in the other AZ.
B. A web tier deployed across 3 AZs with 2 EC2 (Elastic Compute Cloud) instances in each AZ inside an Auto Scaling Group behind an ELB (elastic load balancer) and an application tier deployed across 3 AZs with 2 EC2 instances in each AZ inside an Auto Scaling Group behind an ELB and one RDS (Relational Database Service) Instance deployed with read replicas in the two other AZs.
C. A web tier deployed across 2 AZs with 3 EC2 (Elastic Compute Cloud) instances in each AZ inside an Auto Scaling Group behind an ELB (elastic load balancer) and an application tier deployed across 2 AZs with 3 EC2 instances m each AZ inside an Auto Scaling Group behind an ELS and a Multi-AZ RDS (Relational Database Service) deployment.
D. A web tier deployed across 3 AZs with 2 EC2 (Elastic Compute Cloud) instances in each AZ Inside an Auto Scaling Group behind an ELB (elastic load balancer). And an application tier deployed across 3 AZs with 2 EC2 instances in each AZ inside an Auto Scaling Group behind an ELB and a Multi- AZ RDS (Relational Database services) deployment.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Amazon RDS Multi-AZ Deployments
Amazon RDS Multi-AZ deployments provide enhanced availability and durability for Database (DB) Instances, making them a natural fit for production database workloads. When you provision a Multi-AZ DB Instance, Amazon RDS automatically creates a primary DB Instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). Each AZ runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. In case of an infrastructure failure (for example, instance hardware failure, storage failure, or network

disruption), Amazon RDS performs an automatic failover to the standby, so that you can resume database operations as soon as the failover is complete. Since the endpoint for your DB Instance remains the same after a failover, your application can resume database operation without the need for manual administrative intervention.
Enhanced Durability
Multi-AZ deployments for the MySQL, Oracle, and PostgreSQL engines utilize synchronous physical replication to keep data on the standby up-to-date with the primary. Multi-AZ deployments for the SQL Server engine use synchronous logical replication to achieve the same result, employing SQL Server-native Mirroring technology. Both approaches safeguard your data in the event of a DB Instance failure or loss of an Availability Zone.
If a storage volume on your primary fails in a Multi-AZ deployment, Amazon RDS automatically initiates a failover to the up-to-date standby. Compare this to a Single-AZ deployment: in case of a Single-AZ database failure, a user-initiated point-in-time-restore operation will be required. This operation can take several hours to complete, and any data updates that occurred after the latest restorable time (typically within the last five minutes) will not be available.
Amazon Aurora employs a highly durable, SSD-backed virtualized storage layer purpose-built for database workloads.
Amazon Aurora automatically replicates your volume six ways, across three Availability Zones. Amazon Aurora storage is fault-tolerant, transparently handling the loss of up to two copies of data without affecting database write availability and up to three copies without affecting read availability. Amazon Aurora storage is also self-healing. Data blocks and disks are continuously scanned for errors and replaced automatically.
Increased Availability
You also benefit from enhanced database availability when running Multi-AZ deployments. If an Availability Zone failure or DB Instance failure occurs, your availability impact is limited to the time automatic failover takes to complete: typically under one minute for Amazon Aurora and one to two minutes for other database engines (see the RDS FAQ for details).
The availability benefits of Multi-AZ deployments also extend to planned maintenance and backups. In the case of system upgrades like OS patching or DB Instance scaling, these operations are applied first on the standby, prior to the automatic failover. As a result, your availability impact is, again, only the time required for automatic failover to complete.
Unlike Single-AZ deployments, I/O activity is not suspended on your primary during backup for Multi-AZ deployments for the MySQL, Oracle, and PostgreSQL engines, because the backup is taken from the standby. However, note that you may still experience elevated latencies for a few minutes during backups for Multi-AZ deployments.
On instance failure in Amazon Aurora deployments, Amazon RDS uses RDS Multi-AZ technology to automate failover to one of up to 15 Amazon Aurora Replicas you have created in any of three Availability Zones. If no Amazon Aurora Replicas have been provisioned, in the case of a failure, Amazon RDS will attempt to create a new Amazon Aurora DB instance for you automatically.

QUESTION 45
Your customer wishes to deploy an enterprise application to AWS which will consist of several web servers, several application servers and a small (50GB) Oracle database information is stored, both in the database and the file systems of the various servers. The backup system must support database recovery whole server and whole disk restores, and individual file restores with a recovery time of no more than two hours. They have chosen to use RDS Oracle as the database.
Which backup architecture will meet these requirements?

A. Backup RDS using automated daily DB backups Backup the EC2 instances using AMIs and supplement with file-level backup to S3 using traditional enterprise backup software to provide file level restore
B. Backup RDS using a Multi-AZ Deployment Backup the EC2 instances using Amis, and supplement by copying file system data to S3 to provide file level restore.
C. Backup RDS using automated daily DB backups Backup the EC2 instances using EBS snapshots and supplement with file-level backups to Amazon Glacier using traditional enterprise backup software to provide file level restore
D. Backup RDS database to S3 using Oracle RMAN Backup the EC2 instances using Amis, and supplement with EBS snapshots for individual volume restore.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Point-In-Time Recovery
In addition to the daily automated backup, Amazon RDS archives database change logs. This enables you to recover your database to any point in time during the backup retention period, up to the last five minutes of database usage, Amazon RDS stores multiple copies of your data, but for Single-AZ DB instances these copies are stored in a single availability zone. If for any reason a Single-AZ DB instance becomes unusable, you can use point-in-time recovery to launch a new DB instance with the latest restorable data. For more information on working with point-in-time recovery, go to Restoring a DB Instance to a Specified Time.
Note
Multi-AZ deployments store copies of your data indifferent Availability Zones for greater levels of data durability. For more information on Multi-AZ deployments, see High Availability (Multi-AZ).

QUESTION 46
Your company has HQ in Tokyo and branch offices all over the world and is using a logistics software with a multi-regional deployment on AWS in Japan, Europe and USA. The logistic software has a 3-tier architecture and currently uses MySQL 5.5 for data persistence. Each region has deployed its own database.
In the HQ region you run an hourly batch process reading data from every region to compute cross-regional reports that are sent by email to all offices this batch process must be completed as fast as possible to quickly optimize logistics.
How do you build the database architecture in order to meet the requirements?

A. For each regional deployment, use RDS MySQL with a master in the region and a read replica in the HQ region
B. For each regional deployment, use MySQL on EC2 with a master in the region and send hourly EBS snapshots to the HQ region
C. For each regional deployment, use RDS MySQL with a master in the region and send hourly RDS snapshots to the HQ region
D. For each regional deployment, use MySQL on EC2 with a master in the region and use S3 to copy data files hourly to the HQ region

E. Use Direct Connect to connect all regional MySQL deployments to the HQ region and reduce network latency for the batch process

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 47
A web design company currently runs several FTP servers that their 250 customers use to upload and download large graphic files They wish to move this system to AWS to make it more scalable, but they wish to maintain customer privacy and Keep costs to a minimum.
What AWS architecture would you recommend?

A. ASK their customers to use an S3 client instead of an FTP client. Create a single S3 bucket Create an IAM user for each customer Put the IAM Users in a Group that has an IAM policy that permits access to sub-directories within the bucket via use of the 'username' Policy variable.
B. Create a single S3 bucket with Reduced Redundancy Storage turned on and ask their customers to use an S3 client instead of an FTP client Create a bucket for each customer with a Bucket Policy that permits access only to that one customer.
C. Create an auto-scaling group of FTP servers with a scaling policy to automatically scale-in when minimum network traffic on the auto-scaling group is below a given threshold. Load a central list of ftp users from S3 as part of the user Data startup script on each Instance.
D. Create a single S3 bucket with Requester Pays turned on and ask their customers to use an S3 client instead of an FTP client Create a bucket tor each customer with a Bucket Policy that permits access only to that one customer.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 48
You would like to create a mirror image of your production environment in another region for disaster recovery purposes.
Which of the following AWS resources do not need to be recreated in the second region? (Choose 2 answers)

A. Route 53 Record Sets
B. IAM Roles
C. Elastic IP Addresses (EIP)
D. EC2 Key Pairs
E. Launch configurations
F. Security Groups

Correct Answer: AC
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
http://ltech.com/wp-content/themes/optimize/download/AWS_Disaster_Recovery.pdf (page 6)

QUESTION 49
Your company currently has a 2-tier web application running in an on-premises data center. You have experienced several infrastructure failures in the past two months resulting in significant financial losses. Your CIO is strongly agreeing to move the application to AWS. While working on achieving buy-in from the other company executives, he asks you to develop a disaster recovery plan to help improve Business continuity in the short term. He specifies a target Recovery Time Objective (RTO) of 4 hours and a Recovery Point Objective (RPO) of 1 hour or less. He also asks you to implement the solution within 2 weeks.
Your database is 200GB in size and you have a 20Mbps Internet connection. How would you do this while minimizing costs?

A. Create an EBS backed private AMI which includes a fresh install of your application. Develop a CloudFormation template which includes your AMI and the required EC2, AutoScaling, and ELB resources to support deploying the application across Multiple- Availability-Zones. Asynchronously replicate transactions from your on-premises database to a database instance in AWS across a secure VPN connection.
B. Deploy your application on EC2 instances within an Auto Scaling group across multiple availability zones. Asynchronously replicate transactions from your on-premises database to a database instance in AWS across a secure VPN connection.
C. Create an EBS backed private AMI which includes a fresh install of your application. Setup a script in your data center to backup the local database every 1 hour and to encrypt and copy the resulting file to an S3 bucket using multi-part

upload.

D. Install your application on a compute-optimized EC2 instance capable of supporting the application's average load. Synchronously replicate transactions from your on-premises database to a database instance in AWS across a secure Direct Connect connection.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Overview of Creating Amazon EBS-Backed AMIs
First, launch an instance from an AMI that's similar to the AMI that you'd like to create. You can connect to your instance and customize it. When the instance is configured correctly, ensure data integrity by stopping the instance before you create an AMI, then create the image. When you create an Amazon EBS-backed AMI, we automatically register it for you.
Amazon EC2 powers down the instance before creating the AMI to ensure that everything on the instance is stopped and in a consistent state during the creation process. If you're confident that your instance is in a consistent state appropriate for AMI creation, you can tell Amazon EC2 not to power down and reboot the instance. Some file systems, such as XFS, can freeze and unfreeze activity, making it safe to create the image without rebooting the instance.
During the AMI-creation process, Amazon EC2 creates snapshots of your instance's root volume and any other EBS volumes attached to your instance. If any volumes attached to the instance are encrypted, the new AMI only launches successfully on instances that support Amazon EBS encryption. For more information, see Amazon EBS Encryption.
Depending on the size of the volumes, it can take several minutes for the AMI-creation process to complete (sometimes up to 24 hours). You may find it more efficient to create snapshots of your volumes prior to creating your AMI. This way, only small, incremental snapshots need to be created when the AMI is created, and the process completes more quickly (the total time for snapshot creation remains the same). For more information, see Creating an Amazon EBS Snapshot.
After the process completes, you have a new AMI and snapshot created from the root volume of the instance. When you launch an instance using the new AMI, we create a new EBS volume for its root volume using the snapshot. Both the AMI and the snapshot incur charges to your account until you delete them. For more information, see Deregistering Your AMI.
If you add instance-store volumes or EBS volumes to your instance in addition to the root device volume, the block device mapping for the new AMI contains information for these volumes, and the block device mappings for instances that you launch from the new AMI automatically contain information for these volumes. The instance-store volumes specified in the block device mapping for the new instance are new and don't contain any data from the instance store volumes of the instance you used to create the AMI. The data on EBS volumes persists. For more information, see Block Device Mapping.

QUESTION 50
An enterprise wants to use a third-party SaaS application. The SaaS application needs to have access to issue several API commands to discover Amazon EC2 resources running within the enterprise's account The enterprise has internal security policies that require any outside access to their environment must conform to the principles of least privilege and there must be controls in place to ensure that the credentials used by the SaaS vendor cannot be used by any other third party.
Which of the following would meet all of these conditions?

A. From the AWS Management Console, navigate to the Security Credentials page and retrieve the access and secret key for your account.
B. Create an IAM user within the enterprise account assign a user policy to the IAM user that allows only the actions required by the SaaS application create a new access and secret key for the user and provide these credentials to the SaaS provider.
C. Create an IAM role for cross-account access allows the SaaS provider's account to assume the role and assign it a policy that allows only the actions required by the SaaS application.
D. Create an IAM role for EC2 instances, assign it a policy that allows only the actions required tor the Saas application to work, provide the role ARN to the SaaS provider to use when launching their application instances.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Granting Cross-account Permission to objects It Does Not Own
In this example scenario, you own a bucket and you have enabled other AWS accounts to upload objects. That is, your bucket can have objects that other AWS accounts own.
Now, suppose as a bucket owner, you need to grant cross-account permission on objects, regardless of who the owner is, to a user in another account. For example, that user could be a billing application that needs to access object metadata. There are two core issues:
The bucket owner has no permissions on those objects created by other AWS accounts. So for the bucket owner to grant permissions on objects it does not own, the object owner, the AWS account that created the objects, must first grant permission to the bucket owner. The bucket owner can then delegate those permissions.
Bucket owner account can delegate permissions to users in its own account but it cannot delegate permissions to other AWS accounts, because cross-account delegation is not supported.
In this scenario, the bucket owner can create an AWS Identity and Access Management (IAM) role with permission to access objects, and grant another AWS account permission to assume the role temporarily enabling it to access objects in the bucket.

Background: Cross-Account Permissions and Using IAM Roles
IAM roles enable several scenarios to delegate access to your resources, and cross-account access is one of the key scenarios. In this example, the bucket owner, Account A, uses an IAM role to temporarily delegate object access cross¬account to users in another AWS account, Account C. Each IAM role you create has two policies attached to it:
A trust policy identifying another AWS account that can assume the role.
An access policy defining what permissions—for example, s3:GetObject—are allowed when someone assumes the role. For a list of permissions you can specify in a policy, see Specifying Permissions in a Policy.
The AWS account identified in the trust policy then grants its user permission to assume the role. The user can then do the following to access objects:
Assume the role and, in response, get temporary security credentials.
Using the temporary security credentials, access the objects in the bucket.
For more information about IAM roles, go to Roles (Delegation and Federation) in IAM User Guide. The following is a summary of the walkthrough steps:

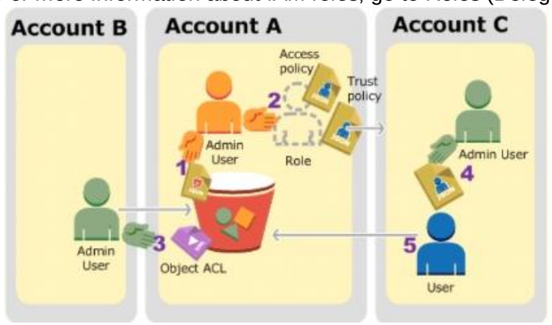

Account A administrator user attaches a bucket policy granting Account B conditional permission to upload objects.
Account A administrator creates an IAM role, establishing trust with Account C, so users in that account can access Account A. The access policy attached to the role limits what user in Account C can do when the user accesses Account A.
Account B administrator uploads an object to the bucket owned by Account A, granting full-control permission to the bucket owner.
Account C administrator creates a user and attaches a user policy that allows the user to assume the role.
User in Account C first assumes the role, which returns the user temporary security credentials. Using those temporary credentials, the user then accesses objects in the bucket.
For this example, you need three accounts. The following table shows how we refer to these accounts and the administrator users in these accounts. Per IAM guidelines (see About Using an Administrator User to Create Resources and Grant Permissions) we do not use the account root credentials in this walkthrough. Instead, you create an administrator user in each account and use those credentials in creating resources and granting them permissions

AWS Account ID	Account Referred To As	Administrator User in the Account
1111-1111-1111	Account A	AccountAadmin
2222-2222-2222	Account B	AccountBadmin
3333-3333-3333	Account C	AccountCadmin

QUESTION 51
A customer has a 10 GB AWS Direct Connect connection to an AWS region where they have a web application hosted on Amazon Elastic Computer Cloud (EC2). The application has dependencies on an on-premises mainframe database that uses a BASE (Basic Available. Sort stale Eventual consistency) rather than an ACID (Atomicity. Consistency isolation. Durability) consistency model. The application is exhibiting undesirable behavior because the database is not able to handle the volume of writes.
How can you reduce the load on your on-premises database resources in the most cost-effective way?

A. Use an Amazon Elastic Map Reduce (EMR) S3DistCp as a synchronization mechanism between the on-premises database and a Hadoop cluster on AWS.
B. Modify the application to write to an Amazon SQS queue and develop a worker process to flush the queue to the on-premises database.
C. Modify the application to use DynamoDB to feed an EMR cluster which uses a map function to write to the on-premises database.
D. Provision an RDS read-replica database on AWS to handle the writes and synchronize the two databases using Data Pipeline.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Reference:
https://aws.amazon.com/blogs/aws/category/amazon-elastic-map-reduce/

QUESTION 52
You are responsible for a legacy web application whose server environment is approaching end of life You would like to migrate this application to AWS as quickly as possible, since the application environment currently has the following limitations:
The VM's single 10GB VMDK is almost full;
Me virtual network interface still uses the 10Mbps driver, which leaves your 100Mbps WAN connection completely underutilized;
It is currently running on a highly customized. Windows VM within a VMware environment;
You do not have me installation media;
This is a mission critical application with an RTO (Recovery Time Objective) of 8 hours. RPO (Recovery Point Objective) of 1 hour.
How could you best migrate this application to AWS while meeting your business continuity requirements?

A. Use the EC2 VM Import Connector for vCenter to import the VM into EC2.
B. Use Import/Export to import the VM as an ESS snapshot and attach to EC2.
C. Use S3 to create a backup of the VM and restore the data into EC2.
D. Use me ec2-bundle-instance API to Import an Image of the VM into EC2

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 53
An AWS customer runs a public blogging website. The site users upload two million blog entries a month. The average blog entry size is 200 KB. The access rate to blog entries drops to negligible 6 months after publication and users rarely access a blog entry 1 year after publication. Additionally, blog entries have a high update rate during the first 3 months following publication, this drops to no updates after 6 months. The customer wants to use CloudFront to improve his user's load times.
Which of the following recommendations would you make to the customer?

A. Duplicate entries into two different buckets and create two separate CloudFront distributions where S3 access is restricted only to Cloud Front identity
B. Create a CloudFront distribution with "US Europe" price class for US/Europe users and a different CloudFront distribution with “All Edge Locations" for the remaining users.
C. Create a CloudFront distribution with S3 access restricted only to the CloudFront identity and partition the blog entry's location in S3 according to the month it was uploaded to be used with CloudFront behaviors.
D. Create a CloudFront distribution with Restrict Viewer Access Forward Query string set to true and minimum TTL of 0.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 54
You are implementing a URL whitelisting system for a company that wants to restrict outbound HTTP'S connections to specific domains from their EC2- hosted applications you deploy a single EC2 instance running proxy software and configure It to accept traffic from all subnets and EC2 instances in the VPC. You configure the proxy to only pass through traffic to domains that you define in its whitelist configuration You have a nightly maintenance window or 10 minutes where all instances fetch new software updates. Each update Is about 200MB In size and there are 500 instances In the VPC that routinely fetch updates After a few days you notice that some machines are failing to successfully download some, but not all of their updates within the maintenance window. The download URLs used for these updates are correctly listed in the proxy's whitelist configuration and you are able to access them manually using a web browser on the instances.
What might be happening? (Choose 2)

A. You are running the proxy on an undersized EC2 instance type so network throughput is not sufficient for all instances to download their updates in time.
B. You are running the proxy on a sufficiently-sized EC2 instance in a private subnet and its network throughput is being throttled by a NAT running on an undersized EC2 instance.
C. The route table for the subnets containing the affected EC2 instances is not configured to direct network traffic for the software update locations to the proxy.
D. You have not allocated enough storage to the EC2 instance running the proxy so the network buffer is filling up, causing some requests to fail.
E. You are running the proxy in a public subnet but have not allocated enough EIPs to support the needed network throughput through the Internet Gateway (IGW).

Correct Answer: AB
Section: (none)
Explanation

Explanation/Reference:

QUESTION 55
Company B is launching a new game app for mobile devices. Users will log into the game using their existing social media account to streamline data capture. Company B would like to directly save player data and scoring information from the mobile app to a DynamoDS table named Score Data When a user saves their game the progress data will be stored to the Game state S3 bucket.
What is the best approach for storing data to DynamoDB and S3?

A. Use an EC2 Instance that is launched with an EC2 role providing access to the Score Data DynamoDB table and the GameState S3 bucket that communicates with the mobile app via web services.
B. Use temporary security credentials that assume a role providing access to the Score Data DynamoDB table and the Game State S3 bucket using web identity federation.
C. Use Login with Amazon allowing users to sign in with an Amazon account providing the mobile app with access to the Score Data DynamoDB table and the Game State S3 bucket.
D. Use an IAM user with access credentials assigned a role providing access to the Score Data DynamoDB table and the Game State S3 bucket for distribution with the mobile app.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Web Identity Federation
Imagine that you are creating a mobile app that accesses AWS resources, such as a game that runs on a mobile device
and stores player and score information using Amazon S3 and DynamoDB.
When you write such an app, you'll make requests to AWS services that must be signed with an AWS access key.
However, we strongly recommend that you do not embed or distribute long-term AWS credentials with apps that a user downloads to a device, even in an encrypted store. Instead, build your app so that it requests temporary AWS security credentials dynamically when needed using web identity federation. The supplied temporary credentials map to an AWS role that has only the permissions needed to perform the tasks required by the mobile app.
With web identity federation, you don't need to create custom sign-in code or manage your own user identities. Instead, users of your app can sign in using a well-known identity provider (IdP) —such as Login with Amazon, Facebook, Google, or any other OpenID Connect (OIDC)-compatible IdP, receive an authentication token, and then exchange that token for temporary security credentials in AWS that map to an IAM role with permissions to use the resources in your AWS
account. Using an IdP helps you keep your AWS account secure, because you don't have to embed and distribute long¬
term security credentials with your application.
For most scenarios, we recommend that you use Amazon Cognito because it acts as an identity broker and does much of the federation work for you. For details, see the following section, Using Amazon Cognito for Mobile Apps.
If you don't use Amazon Cognito, then you must write code that interacts with a web IdP (Login with Amazon, Facebook, Google, or any other OIDC-compatible IdP) and then calls the AssumeRoleWithWebIdentity API to trade the authentication token you get from those IdPs for AWS temporary security credentials. If you have already used this approach for existing apps, you can continue to use it.
Using Amazon Cognito for Mobile Apps
The preferred way to use web identity federation is to use Amazon Cognito. For example, Adele the developer is building a game for a mobile device where user data such as scores and profiles is stored in Amazon S3 and Amazon DynamoDB. Adele could also store this data locally on the device and use Amazon Cognito to keep it synchronized across devices.
She knows that for security and maintenance reasons, long-term AWS security credentials should not be distributed with the game. She also knows that the game might have a large number of users. For all of these reasons, she does not want to create new user identities in IAM for each player. Instead, she builds the game so that users can sign in using an
identity that they've already established with a well-known identity provider, such as Login with Amazon, Facebook,
Google, or any OpenID Connect (OIDC)-compatible identity provider. Her game can take advantage of the authentication mechanism from one of these providers to validate the user's identity.
To enable the mobile app to access her AWS resources, Adele first registers for a developer ID with her chosen IdPs. She also configures the application with each of these providers. In her AWS account that contains the Amazon S3 bucket and DynamoDB table for the game, Adele uses Amazon Cognito to create IAM roles that precisely define permissions that the game needs. If she is using an OIDC IdP, she also creates an IAM OIDC identity provider entity to establish trust between her AWS account and the IdP.
In the app's code, Adele calls the sign-in interface for the IdP that she configured previously. The IdP handles all the
details of letting the user sign in, and the app gets an OAuth access token or OIDC ID token from the provider. Adele's
app can trade this authentication information for a set of temporary security credentials that consist of an AWS access
key ID, a secret access key, and a session token. The app can then use these credentials to access web services offered
by AWS. The app is limited to the permissions that are defined in the role that it assumes.
The following figure shows a simplified flow for how this might work, using Login with Amazon as the IdP. For Step 2, the app can also use Facebook, Google, or any OIDC-compatible identity provider, but that's not shown here.
Sample workflow using Amazon Cognito to federate users for a mobile application

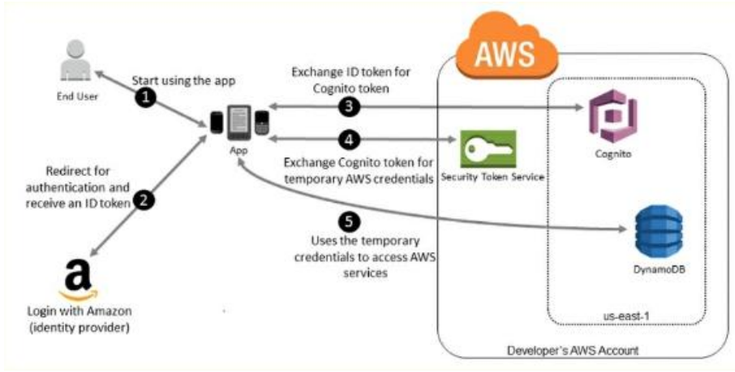

A customer starts your app on a mobile device. The app asks the user to sign in.
The app uses Login with Amazon resources to accept the user's credentials.
The app uses Cognito APIs to exchange the Login with Amazon ID token for a Cognito token.
The app requests temporary security credentials from AWS STS, passing the Cognito token.
The temporary security credentials can be used by the app to access any AWS resources required by the app to operate. The role associated with the temporary security credentials and its assigned policies determines what can be accessed.
Use the following process to configure your app to use Amazon Cognito to authenticate users and give your app access to AWS resources. For specific steps to accomplish this scenario, consult the documentation for Amazon Cognito.
(Optional) Sign up as a developer with Login with Amazon, Facebook, Google, or any other OpenID Connect (OIDC)— compatible identity provider and configure one or more apps with the provider. This step is optional because Amazon Cognito also supports unauthenticated (guest) access for your users.
Go to Amazon Cognito in the AWS Management Console. Use the Amazon Cognito wizard to create an identity pool, which is a container that Amazon Cognito uses to keep end user identities organized for your apps. You can share identity pools between apps. When you set up an identity pool, Amazon Cognito creates one or two IAM roles (one for authenticated identities, and one for unauthenticated "guest" identities) that define permissions for Amazon Cognito users.
Download and integrate the AWS SDK for iOS or the AWS SDK for Android with your app, and import the files required to use Amazon Cognito.
Create an instance of the Amazon Cognito credentials provider, passing the identity pool ID, your AWS account number, and the Amazon Resource Name (ARN) of the roles that you associated with the identity pool. The Amazon Cognito wizard in the AWS Management Console provides sample code to help you get started.
When your app accesses an AWS resource, pass the credentials provider instance to the client object, which passes temporary security credentials to the client. The permissions for the credentials are based on the role or roles that you defined earlier.

QUESTION 56
Your company is getting ready to do a major public announcement of a social media site on AWS. The website is running on EC2 instances deployed across multiple Availability Zones with a Multi-AZ RDS MySQL Extra Large DB Instance. The site performs a high number of small reads and writes per second and relies on an eventual consistency model. After comprehensive tests you discover that there is read contention on RDS MySQL.
Which are the best approaches to meet these requirements? (Choose 2 answers)

A. Deploy ElastiCache in-memory cache running in each availability zone
B. Implement sharding to distribute load to multiple RDS MySQL instances
C. Increase the RDS MySQL Instance size and Implement provisioned IOPS
D. Add an RDS MySQL read replica in each availability zone

Correct Answer: AC
Section: (none)
Explanation

Explanation/Reference:

QUESTION 57
You are designing an intrusion detection prevention (IDS/IPS) solution for a customer web application in a single VPC. You are considering the options for implementing IOS IPS protection for traffic coming from the Internet.
Which of the following options would you consider? (Choose 2 answers)

A. Implement IDS/IPS agents on each Instance running In VPC
B. Configure an instance in each subnet to switch its network interface card to promiscuous mode and analyze network traffic.
C. Implement Elastic Load Balancing with SSL listeners In front of the web applications
D. Implement a reverse proxy layer in front of web servers and configure IDS/IPS agents on each reverse proxy server.

Correct Answer: BD
Section: (none)

Explanation

Explanation/Reference:

QUESTION 58

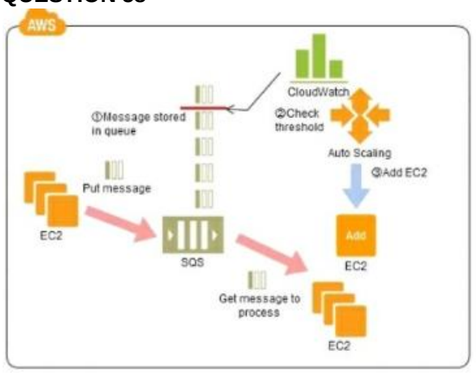

Refer to the architecture diagram above of a batch processing solution using Simple Queue Service (SQS) to set up a message queue between EC2 instances which are used as batch processors Cloud Watch monitors the number of Job requests (queued messages) and an Auto Scaling group adds or deletes batch servers automatically based on parameters set in Cloud Watch alarms.
You can use this architecture to implement which of the following features in a cost effective and efficient manner?

A. Reduce the overall lime for executing jobs through parallel processing by allowing a busy EC2 instance that receives a message to pass it to the next instance in a daisy-chain setup.
B. Implement fault tolerance against EC2 instance failure since messages would remain in SQS and worn can continue with recovery of EC2 instances implement fault tolerance against SQS failure by backing up messages to S3.
C. Implement message passing between EC2 instances within a batch by exchanging messages through SQS.
D. Coordinate number of EC2 instances with number of job requests automatically thus Improving cost effectiveness.
E. Handle high priority jobs before lower priority jobs by assigning a priority metadata field to SQS messages.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
There are cases where a large number of batch jobs may need processing, and where the jobs may need to be re¬prioritized.
For example, one such case is one where there are differences between different levels of services for unpaid users versus subscriber users (such as the time until publication) in services enabling, for example, presentation files to be uploaded for publication from a web browser. When the user uploads a presentation file, the conversion processes, for example, for publication are performed as batch processes on the system side, and the file is published after the conversion. Is it then necessary to be able to assign the level of priority to the batch processes for each type of subscriber?
Explanation of the Cloud Solution/Pattern
A queue is used in controlling batch jobs. The queue need only be provided with priority numbers. Job requests are controlled by the queue, and the job requests in the queue are processed by a batch server. In Cloud computing, a highly reliable queue is provided as a service, which you can use to structure a highly reliable batch system with ease. You may prepare multiple queues depending on priority levels, with job requests put into the queues depending on their priority levels, to apply prioritization to batch processes. The performance (number) of batch servers corresponding to a queue must be in accordance with the priority level thereof.
Implementation
In AWS, the queue service is the Simple Queue Service (SQS). Multiple SQS queues may be prepared to prepare queues for individual priority levels (with a priority queue and a secondary queue). Moreover, you may also use the message Delayed Send function to delay process execution.
Use SQS to prepare multiple queues for the individual priority levels.
Place those processes to be executed immediately (job requests) in the high priority queue.
Prepare numbers of batch servers, for processing the job requests of the queues, depending on the priority levels. Queues have a message "Delayed Send" function. You can use this to delay the time for starting a process.
Configuration

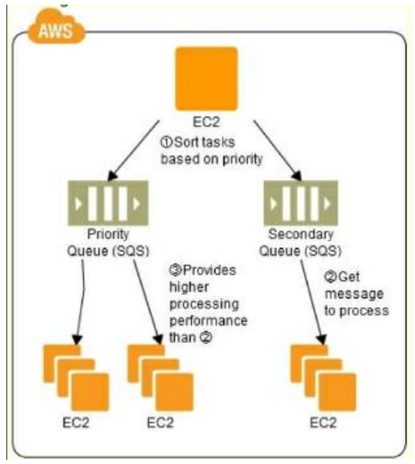

Benefits
You can increase or decrease the number of servers for processing jobs to change automatically the processing speeds of the priority queues and secondary queues.
You can handle performance and service requirements through merely increasing or decreasing the number of EC2 instances used in job processing.
Even if an EC2 were to fail, the messages (jobs) would remain in the queue service, enabling processing to be continued immediately upon recovery of the EC2 instance, producing a system that is robust to failure.
Cautions
Depending on the balance between the number of EC2 instances for performing the processes and the number of messages that are queued, there may be cases where processing in the secondary queue may be completed first, so you need to monitor the processing speeds in the primary queue and the secondary queue.

QUESTION 59
An International company has deployed a multi-tier web application that relies on DynamoDB in a single region. For regulatory reasons they need disaster recovery capability In a separate region with a Recovery Time Objective of 2 hours and a Recovery Point Objective of 24 hours. They should synchronize their data on a regular basis and be able to provision me web application rapidly using CloudFormation.
The objective is to minimize changes to the existing web application, control the throughput of DynamoDB used for the synchronization of data and synchronize only the modified elements.
Which design would you choose to meet these requirements?

A. Use AWS data Pipeline to schedule a DynamoDB cross region copy once a day, create a "Lastupdated" attribute in your DynamoDB table that would represent the timestamp of the last update and use it as a filter.
B. Use EMR and write a custom script to retrieve data from DynamoDB in the current region using a SCAN operation and push it to DynamoDB in the second region.
C. Use AWS data Pipeline to schedule an export of the DynamoDB table to S3 in the current region once a day then schedule another task immediately after it that will import data from S3 to DynamoDB in the other region.
D. Send also each Ante into an SQS queue in me second region; use an auto-scaling group behind the SQS queue to replay the write in the second region.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 60
You are designing a social media site and are considering how to mitigate distributed denial-of-service (DDoS) attacks.
Which of the below are viable mitigation techniques? (Choose 3)

A. Add multiple elastic network interfaces (ENIs) to each EC2 instance to increase the network bandwidth.
B. Use dedicated instances to ensure that each instance has the maximum performance possible.
C. Use an Amazon CloudFront distribution for both static and dynamic content.
D. Use an Elastic Load Balancer with auto scaling groups at the web. App and Amazon Relational Database Service (RDS) tiers
E. Add alert Amazon CloudWatch to look for high Network in and CPU utilization.
F. Create processes and capabilities to quickly add and remove rules to the instance OS firewall.

Correct Answer: CEF
Section: (none)
Explanation

Explanation/Reference:

QUESTION 61
You must architect the migration of a web application to AWS. The application consists of Linux web servers running a custom web server. You are required to save the logs generated from the application to a durable location.
What options could you select to migrate the application to AWS? (Choose 2)

A. Create an AWS Elastic Beanstalk application using the custom web server platform. Specify the web server executable and the application project and source files. Enable log file rotation to Amazon Simple Storage Service (S3).
B. Create Dockerfile for the application. Create an AWS OpsWorks stack consisting of a custom layer. Create custom recipes to install Docker and to deploy your Docker container using the Dockerfile. Create customer recipes to install and configure the application to publish the logs to Amazon CloudWatch Logs.
C. Create Dockerfile for the application. Create an AWS OpsWorks stack consisting of a Docker layer that uses the Dockerfile. Create custom recipes to install and configure Amazon Kineses to publish the logs into Amazon CloudWatch.
D. Create a Dockerfile for the application. Create an AWS Elastic Beanstalk application using the Docker platform and the Dockerfile. Enable logging the Docker configuration to automatically publish the application logs. Enable log file rotation to Amazon S3.
E. Use VM import/Export to import a virtual machine image of the server into AWS as an AMI. Create an Amazon Elastic Compute Cloud (EC2) instance from AMI, and install and configure the Amazon CloudWatch Logs agent. Create a new AMI from the instance. Create an AWS Elastic Beanstalk application using the AMI platform and the new AMI.

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 62
A web company is looking to implement an external payment service into their highly available application deployed in a VPC Their application EC2 instances are behind a public lacing ELB Auto scaling is used to add additional instances as traffic increases under normal load the application runs 2 instances in the Auto Scaling group but at peak it can scale 3x in size. The application instances need to communicate with the payment service over the Internet which requires whitelisting of all public IP addresses used to communicate with it. A maximum of 4 whitelisting IP addresses are allowed at a time and can be added through an API.
How should they architect their solution?

A. Route payment requests through two NAT instances setup for High Availability and whitelist the Elastic IP addresses attached to the MAT instances
B. Whitelist the VPC Internet Gateway Public IP and route payment requests through the Internet Gateway.
C. Whitelist the ELB IP addresses and route payment requests from the Application servers through the ELB.
D. Automatically assign public IP addresses to the application instances in the Auto Scaling group and run a script on boot that adds each instances public IP address to the payment validation whitelist API.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 63
Your website is serving on-demand training videos to your workforce. Videos are uploaded monthly in high resolution MP4 format. Your workforce is distributed globally often on the move and using company-provided tablets that require the HTTP Live Streaming (HLS) protocol to watch a video. Your company has no video transcoding expertise and it required you may need to pay for a consultant.
How do you implement the most cost-efficient architecture without compromising high availability and quality of video delivery ?

A. A video transcoding pipeline running on EC2 using SQS to distribute tasks and Auto Scaling to adjust the number of nodes depending on the length of the queue. EBS volumes to host videos and EBS snapshots to incrementally backup original files after a few days. CloudFront to serve HLS transcoded videos from EC2.
B. Elastic Transcoder to transcode original high-resolution MP4 videos to HLS. EBS volumes to host videos and EBS snapshots to incrementally backup original files after a few days. CloudFront to serve HLS transcoded videos from EC2.
C. Elastic Transcoder to transcode original high-resolution MP4 videos to HLS. S3 to host videos with Lifecycle Management to archive original files to Glacier after a few days. CloudFront to serve HLS transcoded videos from S3.
D. A video transcoding pipeline running on EC2 using SQS to distribute tasks and Auto Scaling to adjust the number of nodes depending on the length of the queue. S3 to host videos with Lifecycle Management to archive all files to Glacier after a few days. CloudFront to serve HLS transcoded videos from Glacier.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 64
A customer has established an AWS Direct Connect connection to AWS. The link is up and routes are being advertised from the customer's end, however the customer is unable to connect from EC2 instances inside its VPC to servers residing in its datacenter.
Which of the following options provide a viable solution to remedy this situation? (Choose 2)

A. Add a route to the route table with an iPsec VPN connection as the target.
B. Enable route propagation to the virtual pinnate gateway (VGW).
C. Enable route propagation to the customer gateway (CGW).
D. Modify the route table of all Instances using the 'route' command.
E. Modify the Instances VPC subnet route table by adding a route back to the customer's on-premises environment.

Correct Answer: AC
Section: (none)
Explanation

Explanation/Reference:

QUESTION 65
You are running a news website in the eu-west-1 region that updates every 15 minutes. The website has a world-wide audience it uses an Auto Scaling group behind an Elastic Load Balancer and an Amazon RDS database Static content resides on Amazon S3, and is distributed through Amazon CloudFront. Your Auto Scaling group is set to trigger a scale up event at 60% CPU utilization, you use an Amazon RDS extra large DB instance with 10.000 Provisioned IOPS its CPU utilization is around 80%. While freeable memory is in the 2 GB range.
Web analytics reports show that the average load time of your web pages is around 1.5 to 2 seconds, but your SEO consultant wants to bring down the average load time to under 0.5 seconds.
How would you improve page load times for your users? (Choose 3 answers)

A. Lower the scale up trigger of your Auto Scaling group to 30% so it scales more aggressively.
B. Add an Amazon ElastiCache caching layer to your application for storing sessions and frequent DB queries
C. Configure Amazon CloudFront dynamic content support to enable caching of re-usable content from your site
D. Switch the Amazon RDS database to the high memory extra large Instance type
E. Set up a second installation in another region, and use the Amazon Route 53 latency-based routing feature to select the right region.

Correct Answer: ABD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 66
A corporate web application is deployed within an Amazon Virtual Private Cloud (VPC) and is connected to the corporate data center via an IPSec VPN. The application must authenticate against the on-premises LDAP server. After authentication, each logged-in user can only access an Amazon Simple Storage Space (S3) keyspace specific to that user.
Which two approaches can satisfy these objectives? (Choose 2)

A. Develop an identity broker that authenticates against I AM security Token service to assume a I AM 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.
B. 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 to assume that IAM role. The application can use the temporary credentials to access the appropriate S3 bucket.
C. Develop an identity broker that authenticates against LDAP and then calls IAM Security Token Service 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.
D. 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.
E. The application authenticates against IAM Security Token Service using the LDAP credentials the application uses those temporary AWS security credentials to access the appropriate S3 bucket.

Correct Answer: BC
Section: (none)
Explanation

Explanation/Reference:

QUESTION 67
Your company previously configured a heavily used, dynamically routed VPN connection between your on-premises data center and AWS. You recently provisioned a DirectConnect connection and would like to start using the new connection.
After configuring DirectConnect settings in the AWS Console, which of the following options win provide the most seamless transition for your users?

A. Delete your existing VPN connection to avoid routing loops configure your DirectConnect router with the appropriate settings and verity network traffic is leveraging DirectConnect.
B. Configure your DirectConnect router with a higher BGP priority man your VPN router, verify network traffic is leveraging DirectConnect and then delete your existing VPN connection.
C. Update your VPC route tables to point to the DirectConnect connection configure your DirectConnect router with the appropriate settings verify network traffic is leveraging DirectConnect and then delete the VPN connection.
D. Configure your DirectConnect router, update your VPC route tables to point to the DirectConnect connection, configure your VPN connection with a higher BGP priority. And verify network traffic is leveraging the DirectConnect connection.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 68
Your company hosts a social media website for storing and sharing documents. The web application allows user to upload large files while resuming and pausing the upload as needed. Currently, files are uploaded to your PHP front end backed by Elastic load Balancing and an autoscaling fleet of Amazon Elastic Compute Cloud (EC2) instances that scale upon average of bytes received (Networking). After a file has been uploaded, it is copied to Amazon Simple Storage Service (S3). Amazon EC2 instances use an AWS Identity and Access Management (IAM) role that allows Amazon S3 uploads. Over the last six months, your user base and scale have increased significantly, forcing you to increase the Auto Scaling group's Max parameter a few times. Your CFO is concerned about rising costs and has asked you to adjust the architecture where needed to better optimize costs.
Which architecture change could you introduce to reduce costs and still keep your web application secure and scalable?

A. Replace the Auto Scaling launch configuration to include c3.8xlarge instances; those instances can potentially yield a network throuthput of lOgbps.
B. Re-architect your ingest pattern, have the app authenticate against your identity provider, and use your identity provider as a broker fetching temporary AWS credentials from AWS Secure Token Service (GetFederationToken). Securely pass the credentials and S3 endpoint/prefix to your app. Implement client-side logic to directly upload the file to Amazon S3 using the given credentials and S3 prefix.
C. Re-architect your ingest pattern, and move your web application instances into a VPC public subnet. Attach a public IP address for each EC2 instance (using the Auto Scaling launch configuration settings). Use Amazon Route 53 Round Robin records set and HTTP health check to DNS load balance the app requests; this approach will significantly reduce the cost by bypassing Elastic Load Balancing.
D. Re-architect your ingest pattern, have the app authenticate against your identity provider, and use your identity provider as a broker fetching temporary AWS credentials from AWS Secure Token Service (GetFederationToken). Securely pass the credentials and S3 endpoint/prefix to your app. Implement client-side logic that used the S3 multipart upload API to directly upload the file to Amazon S3 using the given credentials and S3 prefix.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 69
You require the ability to analyze a customer's clickstream data on a website so they can do behavioral analysis. Your customer needs to know what sequence of pages and ads their customer clicked on. This data will be used in real time to modify the page layouts as customers click through the site to increase stickiness and advertising click-through.
Which option meets the requirements for captioning and analyzing this data?

A. Log clicks in weblogs by URL store to Amazon S3, and then analyze with Elastic MapReduce
B. Push web clicks by session to Amazon Kinesis and analyze behavior using Kinesis workers
C. Write click events directly to Amazon Redshift and then analyze with SQL
D. Publish web clicks by session to an Amazon SQS queue then periodically drain these events to Amazon RDS and analyze with SQL.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
http://www. slideshare.net/AmazonWebServices/aws-webcast-introduction-to-amazon-kinesis

QUESTION 70
You have deployed a three-tier web application in a VPC with a Cl DR block of 10.0.0.0/28. You initially deploy two web servers, two application servers, two database servers and one NAT instance tor a total of seven EC2 instances. The web. Application and database servers are deployed across two availability zones (AZs). You also deploy an ELB in front of the two web servers, and use Route53 for DNS Web (raffle gradually increases in the first few days following the deployment, so you attempt to double the number of instances in each tier of the application to handle the new load unfortunately some of these new instances fail to launch.
Which of the following could be the root caused? (Choose 2 answers)

A. AWS reserves the first and the last private IP address in each subnet's CIDR block so you do not have enough addresses left to launch all of the new EC2 instances
B. The Internet Gateway (IGW) of your VPC has scaled-up, adding more instances to handle the traffic spike, reducing the number of available private IP addresses for new instance launches
C. The ELB has scaled-up, adding more instances to handle the traffic spike, reducing the number of available private IP addresses for new instance launches
D. AWS reserves one IP address in each subnet's CIDR block for Route53 so you do not have enough addresses left to launch all of the new EC2 instances
E. AWS reserves the first four and the last IP address in each subnet's CIDR block so you do not have enough addresses left to launch all of the new EC2 instances

Correct Answer: CE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 71
Your company produces customer commissioned one-of-a-kind skiing helmets combining nigh fashion with custom technical enhancements Customers can show off their Individuality on the ski slopes and have access to head-up-displays. GPS rear-view cams and any other technical innovation they wish to embed in the helmet.
The current manufacturing process is data rich and complex including assessments to ensure that the custom electronics and materials used to assemble the helmets are to the highest standards Assessments are a mixture of human and automated assessments you need to add a new set of assessment to model the failure modes of the custom electronics using GPUs with CUDA. across a cluster of servers with low latency networking.
What architecture would allow you to automate the existing process using a hybrid approach and ensure that the architecture can support the evolution of processes over time?

A. Use AWS Data Pipeline to manage movement of data & meta-data and assessments Use an auto-scaling group of G2 instances in a placement group.
B. Use Amazon Simple Workflow (SWF) to manages assessments, movement of data & meta-data Use an auto-scaling group of G2 instances in a placement group.
C. Use Amazon Simple Workflow (SWF) to manages assessments movement of data & meta-data Use an auto-scaling group of C3 instances with SR- IOV (Single Root I/O Virtualization).
D. Use AWS data Pipeline to manage movement of data & meta data and assessments use auto-scaling group of C3 with SR-IOV (Single Root I/O virtualization).

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 72
You are designing an SSUTLS solution that requires HTTPS clients to be authenticated by the Web server using client certificate authentication. The solution must be resilient.
Which of the following options would you consider for configuring the web server infrastructure? (Choose 2)

A. Configure ELB with TCP listeners on TCP/443. And place the Web servers behind it.
B. Configure your Web servers with EIPs. Place the Web servers in a Route53 Record Set and configure health checks against all Web servers.
C. Configure ELB with HTTPS listeners, and place the Web servers behind it.
D. Configure your web servers as the origins for a CloudFront distribution. Use custom SSL certificates on your CloudFront distribution.

Correct Answer: AB
Section: (none)
Explanation

Explanation/Reference:

QUESTION 73
You are migrating a legacy client-server application to AWS. The application responds to a specific DNS domain (e g. www.example.com) and has a 2-tier architecture, with multiple application servers and a database server. Remote clients use TCP to connect to the application servers. The application servers need to know the IP address of the clients in order to function properly and are currently taking that information from the TCP socket. A Multi-AZ RDS MySQL instance will be used for the database.
During the migration you can change the application code, but you have to file a change request.

How would you implement the architecture on AWS in order to maximize scalability and high availability?

A. File a change request to implement Alias Resource support in the application. Use Route 53 Alias Resource Record to distribute load on two application servers in different Azs.
B. File a change request to implement Latency Based Routing support in the application. Use Route 53 with Latency Based Routing enabled to distribute load on two application servers in different Azs.
C. File a change request to implement Cross-Zone support in the application. Use an ELB with a TCP Listener and Cross-Zone Load Balancing enabled, two application servers in different AZs.
D. File a change request to implement Proxy Protocol support in the application. Use an ELB with a TCP Listener and Proxy Protocol enabled to distribute load on two application servers in different Azs.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 74
You are designing a personal document-archiving solution for your global enterprise with thousands of employee. Each employee has potentially gigabytes of data to be backed up in this archiving solution. The solution will be exposed to the employees as an application, where they can just drag and drop their files to the archiving system. Employees can retrieve their archives through a web interface. The corporate network has high bandwidth AWS Direct Connect connectivity to AWS.
You have a regulatory requirement that all data needs to be encrypted before being uploaded to the cloud.
How do you implement this in a highly available and cost-efficient way?

A. Manage encryption keys on-premises in an encrypted relational database. Set up an on-premises server with sufficient storage to temporarily store files, and then upload them to Amazon S3, providing a client-side master key.
B. Mange encryption keys in a Hardware Security Module (HSM) appliance on-premises serve r with sufficient storage to temporarily store, encrypt, and upload files directly into Amazon Glacier.
C. Manage encryption keys in Amazon Key Management Service (KMS), upload to Amazon Simple Storage Service (S3) with client-side encryption using a KMS customer master key ID, and configure Amazon S3 lifecycle policies to store each object using the Amazon Glacier storage tier.
D. Manage encryption keys in an AWS CloudHSM appliance. Encrypt files prior to uploading on the employee desktop, and then upload directly into Amazon Glacier.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 75
A company is building a voting system for a popular TV show, viewers win watch the performances then visit the show's website to vote for their favorite performer. It is expected that in a short period of time after the show has finished the site will receive millions of visitors. The visitors will first login to the site using their Amazon.com credentials and then submit their vote. After the voting is completed the page will display the vote totals. The company needs to build the site such that can handle the rapid influx of traffic while maintaining good performance but also wants to keep costs to a minimum.
Which of the design patterns below should they use?

A. Use CloudFront and an Elastic Load balancer in front of an auto-scaled set of web servers, the web servers will first call the Login With Amazon service to authenticate the user then process the users vote and store the result into a multi-AZ Relational Database Service instance.
B. Use CloudFront and the static website hosting feature of S3 with the Javascript SDK to call the Login With Amazon service to authenticate the user, use IAM Roles to gain permissions to a DynamoDB table to store the users vote.
C. Use CloudFront and an Elastic Load Balancer in front of an auto-scaled set of web servers, the web servers will first call the Login with Amazon service to authenticate the user, the web servers will process the users vote and store the result into a DynamoDB table using IAM Roles for EC2 instances to gain permissions to the DynamoDB table.
D. Use CloudFront and an Elastic Load Balancer in front of an auto-scaled set of web servers, the web servers will first call the Login With Amazon service to authenticate the user, the web servers win process the users vote and store the result into an SQS queue using IAM Roles for EC2 Instances to gain permissions to the SQS queue. A set of application servers will then retrieve the items from the queue and store the result into a DynamoDB table.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 76
You are designing a connectivity solution between on-premises infrastructure and Amazon VPC. Your servers on-premises will be communicating with your VPC instances. You will be establishing IPSec tunnels over the Internet You will be using VPN gateways, and terminating the IPSec tunnels on AWS supported customer gateways.
Which of the following objectives would you achieve by implementing an IPSec tunnel as outlined above? Choose 4 answers

A. End-to-end protection of data in transit
B. End-to-end Identity authentication
C. Data encryption across the Internet
D. Protection of data in transit over the Internet
E. Peer identity authentication between VPN gateway and customer gateway
F. Data integrity protection across the Internet

Correct Answer: CDEF
Section: (none)
Explanation

Explanation/Reference:

QUESTION 77
You are responsible for a web application that consists of an Elastic Load Balancing (ELB) load balancer in front of an Auto Scaling group of Amazon Elastic Compute Cloud (EC2) instances. For a recent deployment of a new version of the application, a new Amazon Machine Image (AMI) was created, and the Auto Scaling group was updated with a new launch configuration that refers to this new AMI. During the deployment, you received complaints from users that the website was responding with errors. All instances passed the ELB health checks.
What should you do in order to avoid errors for future deployments? (Choose 2)

A. Add an Elastic Load Balancing health check to the Auto Scaling group. Set a short period for the health checks to operate as soon as possible in order to prevent premature registration of the instance to the load balancer.
B. Enable EC2 instance CloudWatch alerts to change the launch configuration's AMI to the previous one. Gradually terminate instances that are using the new AMI.
C. Set the Elastic Load Balancing health check configuration to target a part of the application that fully tests application health and returns an error if the tests fail.
D. Create a new launch configuration that refers to the new AMI, and associate it with the group. Double the size of the group, wait for the new instances to become healthy, and reduce back to the original size. If new instances do not become healthy, associate the previous launch configuration.
E. Increase the Elastic Load Balancing Unhealthy Threshold to a higher value to prevent an unhealthy instance from going into service behind the load balancer.

Correct Answer: CD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 78
Which is a valid Amazon Resource name (ARN) for I AM?

A. aws:iam::123456789012:instance-profile/Webserver
B. arn:aws:iam::123456789012:instance-profile/Webserver
C. 123456789012:aws:iam::instance-profile/Webserver
D. arn:aws:iam::123456789012::instance-profile/Webserver

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
IAM ARNs
Most resources have a friendly name (for example, a user named Bob or a group named Developers). However, the access policy language requires you to specify the resource or resources using the following Amazon Resource Name (ARN) format.
arn:aws:service:region:account:resource
Where:
service identifies the AWS product. For IAM resources, this is always iam. region is the region the resource resides in. For IAM resources, this is always left blank, account is the AWS account ID with no hyphens (for example, 123456789012). resource is the portion that identifies the specific resource by name.
You can use ARNs in IAM for users (IAM and federated), groups, roles, policies, instance profiles, virtual MFA devices, and server certificates. The following table shows the ARN format for each and an example. The region portion of the ARN is blank because IAM resources are global.

QUESTION 79

Dave is the main administrator in Example Corp., and he decides to use paths to help delineate the users in the company and set up a separate administrator group for each path-based division. Following is a subset of the full list of paths he plans to use:
•/marketing
•/sales
•/legal
Dave creates an administrator group for the marketing part of the company and calls it Marketing_Admin.
He assigns it the /marketing path. The group's ARN is arn:aws:iam::123456789012:group/marketing/Marketing_Admin.
Dave assigns the following policy to the Marketing_Admin group that gives the group permission to use all IAM actions with all groups and users in the / marketing path. The policy also gives the Marketing_Admin group permission to perform any AWS S3 actions on the objects in the portion of the corporate bucket.

```
{
"Version": "2012-10-17",
"Statement": [
{
"Effect": "Deny",
"Action": "iam:*",
"Resource": [
"arn:aws:iam::123456789012:group/marketing/*",
"arn:aws:iam::123456789012:user/marketing/*"
]
},
{
"Effect": "Allow",
"Action": "s3:*",
"Resource": "arn:aws:s3:::example_bucket/marketing/*"
},
{
"Effect": "Allow",
"Action": "s3:ListBucket*",
"Resource": "arn:aws:s3:::example_bucket",
"Condition":{"StringLike":{"s3:prefix": "marketing/*"}}
}
]
}
```

A. True
B. False

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 80
Your fortune 500 company has under taken a TCO analysis evaluating the use of Amazon S3 versus acquiring more hardware The outcome was that ail employees would be granted access to use Amazon S3 for storage of their personal documents.
Which of the following will you need to consider so you can set up a solution that incorporates single sign-on from your corporate AD or LDAP directory and restricts access for each user to a designated user folder in a bucket? (Choose 3)

A. Setting up a federation proxy or identity provider
B. Using AWS Security Token Service to generate temporary tokens
C. Tagging each folder in the bucket
D. Configuring IAM role
E. Setting up a matching IAM user for every user in your corporate directory that needs access to a folder in the bucket

Correct Answer: ABD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 81
A company is running a batch analysis every hour on their main transactional DB, running on an RDS MySQL instance, to populate their central Data Warehouse running on Redshift. During the execution of the batch, their transactional applications are very slow. When the batch completes they need to update the top management dashboard with the new data. The dashboard is produced by another system running on-premises that is currently started when a manually-sent email notifies that an update is required. The on-premises system cannot be modified because is managed by another team.
How would you optimize this scenario to solve performance issues and automate the process as much as possible?

A. Replace RDS with Redshift for the batch analysis and SNS to notify the on-premises system to update the dashboard
B. Replace RDS with Redshift for the oaten analysis and SQS to send a message to the on-premises system to update the dashboard
C. Create an RDS Read Replica for the batch analysis and SNS to notify me on-premises system to update the dashboard
D. Create an RDS Read Replica for the batch analysis and SQS to send a message to the on-premises system to update the dashboard.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 82
You are running a successful multitier web application on AWS and your marketing department has asked you to add a reporting tier to the application. The reporting tier will aggregate and publish status reports every 30 minutes from user-generated information that is being stored in your web application s database. You are currently running a Multi-AZ RDS MySQL instance for the database tier. You also have implemented Elasticache as a database caching layer between the application tier and database tier.
Please select the answer that will allow you to successfully implement the reporting tier with as little impact as possible to your database.

A. Continually send transaction logs from your master database to an S3 bucket and generate the reports off the S3 bucket using S3 byte range requests.
B. Generate the reports by querying the synchronously replicated standby RDS MySQL instance maintained through Multi-AZ.
C. Launch a RDS Read Replica connected to your Multi AZ master database and generate reports by querying the Read Replica.
D. Generate the reports by querying the ElastiCache database caching tier.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Amazon RDS allows you to use read replicas with Multi-AZ deployments. In Multi-AZ deployments for MySQL, Oracle, SQL Server, and PostgreSQL, the data in your primary DB Instance is synchronously replicated to to a standby instance in a different Availability Zone (AZ). Because of their synchronous replication, Multi-AZ deployments for these engines offer greater data durability benefits than do read replicas. (In all Amazon RDS for Aurora deployments, your data is automatically replicated across 3 Availability Zones.)
You can use Multi-AZ deployments and read replicas in conjunction to enjoy the complementary benefits of each. You can simply specify that a given Multi-AZ deployment is the source DB Instance for your Read replicas. That way you gain both the data durability and availability benefits of Multi-AZ deployments and the read scaling benefits of read replicas.
Note that for Multi-AZ deployments, you have the option to create your read replica in an AZ other than that of the primary and the standby for even more redundancy. You can identify the AZ corresponding to your standby by looking at the "Secondary Zone" field of your DB Instance in the AWS Management Console.

QUESTION 83
You are designing a data leak prevention solution for your VPC environment. You want your VPC Instances to be able to access software depots and distributions on the Internet for product updates. The depots and distributions are accessible via third party CDNs by their URLs.
You want to explicitly deny any other outbound connections from your VPC instances to hosts on the internet.
Which of the following options would you consider?

A. Configure a web proxy server in your VPC and enforce URL-based rules for outbound access Remove default routes.
B. Implement security groups and configure outbound rules to only permit traffic to software depots.
C. Move all your instances into private VPC subnets remove default routes from all routing tables and add specific routes to the software depots and distributions only.
D. Implement network access control lists to all specific destinations, with an Implicit deny as a rule.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 84
You have an application running on an EC2 instance which will allow users to download files from a private S3 bucket using a pre-signed URL. Before generating the URL, the application should verify the existence of the file in S3.
How should the application use AWS credentials to access the S3 bucket securely?

A. Use the AWS account access keys; the application retrieves the credentials from the source code of the application.
B. Create an IAM role for EC2 that allows list access to objects In the S3 bucket; launch the Instance with the role, and retrieve the role's credentials from the EC2 instance metadata.
C. Create an IAM user for the application with permissions that allow list access to the S3 bucket; the application retrieves the IAM user credentials from a temporary directory with permissions that allow read access only to the Application user.
D. Create an IAM user for the application with permissions that allow list access to the S3 bucket; launch the instance as the IAM user, and retrieve the IAM user's credentials from the EC2 instance user data.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 85
Your system recently experienced down time during the troubleshooting process. You found that a new administrator mistakenly terminated several production EC2 instances.
Which of the following strategies will help prevent a similar situation in the future?
The administrator still must be able to:
launch, start stop, and terminate development resources.
launch and start production instances.

A. Create an IAM user, which is not allowed to terminate instances by leveraging production EC2 termination protection.
B. Leverage resource based tagging, along with an IAM user which can prevent specific users from terminating production, EC2 resources.
C. Leverage EC2 termination protection and multi-factor authentication, which together require users to authenticate before terminating EC2 instances
D. Create an IAM user and apply an IAM role which prevents users from terminating production EC2 instances.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Working with volumes
When an API action requires a caller to specify multiple resources, you must create a policy statement that allows users to access all required resources. If you need to use a Condition element with one or more of these resources, you must create multiple statements as shown in this example.
The following policy allows users to attach volumes with the tag "volume_user=iam-user-name" to instances with the
tag "department=dev", and to detach those volumes from those instances. If you attach this policy to an IAM group, the aws:username policy variable gives each IAM user in the group permission to attach or detach volumes from the instances with a tag named volume_user that has his or her IAM user name as a value.

```
{
"Version": "2012-10-17",
"Statement": [{
"Effect": "Allow",
"Action": [
"ec2:AttachVolume",
"ec2:DetachVolume"
],
"Resource": "arn:aws:ec2:us-east-1:123456789012:instance/*",
"Condition": {
"StringEquals": {
"ec2:ResourceTag/department": "dev"
}
}
},
{
"Effect": "Allow",
"Action": [
"ec2:AttachVolume",
```

```
"ec2:DetachVolume"
],
"Resource": "arn:aws:ec2:us-east-1:123456789012:volume/*", "Condition": {
"StringEquals": {
"ec2:ResourceTag/volume_user": "${aws:username}"
}
}
}
]
}
```

Launching instances (RunInstances)

The RunInstances API action launches one or more instances. RunInstances requires an AMI and creates an instance; and users can specify a key pair and security group in the request. Launching into EC2-VPC requires a subnet, and creates a network interface. Launching from an Amazon EBS-backed AMI creates a volume. Therefore, the user must have permission to use these Amazon EC2 resources. The caller can also configure the instance using optional parameters to RunInstances, such as the instance type and a subnet. You can create a policy statement that requires users to specify an optional parameter, or restricts users to particular values for a parameter. The examples in this section demonstrate some of the many possible ways that you can control the configuration of an instance that a user can launch.

Note that by default, users don't have permission to describe, start, stop, or terminate the resulting instances. One way to grant the users permission to manage the resulting instances is to create a specific tag for each instance, and then create a statement that enables them to manage instances with that tag. For more information, see 2: Working with instances.

a.AMI

The following policy allows users to launch instances using only the AMIs that have the specified tag, "department=dev", associated with them. The users can't launch instances using other AMIs because the Condition element of the first statement requires that users specify an AMI that has this tag. The users also can't launch into a subnet, as the policy does not grant permissions for the subnet and network interface resources. They can, however, launch into EC2-Classic. The second statement uses a wildcard to enable users to create instance resources, and requires users to specify the key pair project_keypair and the security group sg-1a2b3c4d. Users are still able to launch instances without a key pair.

```
{
Version": "2012-10-17", Statement": [{
"Effect": "Allow",
"Action": "ec2:RunInstances",
"Resource": [
"arn:aws:ec2:region::image/ami-*"
],
"Condition": {
"StringEquals": {
"ec2:ResourceTag/department": "dev"
}
}
},
{
"Effect": "Allow",
"Action": "ec2:RunInstances",
"Resource": [
"arn:aws:ec2:region:account:instance/*",
"arn:aws:ec2:region:account:volume/*",
"arn:aws:ec2:region:account:key-pair/project_keypair",
"arn:aws:ec2:region:account:security-group/sg-1a2b3c4d"
]
}
]
}
```

Alternatively, the following policy allows users to launch instances using only the specified AMIs, ami-9e1670f7 and ami- 45cf5c3c. The users can't launch an instance using other AMIs (unless another statement grants the users permission to do so), and the users can't launch an instance into a subnet.

```
{
"Version": "2012-10-17",
"Statement": [{
"Effect": "Allow",
"Action": "ec2:RunInstances",
"Resource": [
"arn:aws:ec2:region::image/ami-9e1670f7",
"arn:aws:ec2:region::image/ami-45cf5c3c",
"arn:aws:ec2:region:account:instance/*",
"arn:aws:ec2:region:account:volume/*",
"arn:aws:ec2:region:account:key-pair/*",
"arn:aws:ec2:region:account:security-group/*"
]
}
]
```

```
}
```
Alternatively, the following policy allows users to launch instances from all AMIs owned by Amazon. The Condition element of the first statement tests whether ec2:Owner is amazon. The users can't launch an instance using other AMIs (unless another statement grants the users permission to do so). The users are able to launch an instance into a subnet.
```
{
"Version": "2012-10-17",
"Statement": [{
"Effect": "Allow",
"Action": "ec2:RunInstances",
"Resource": [
"arn:aws:ec2:region::image/ami-*"
],
"Condition": {
"StringEquals": {
"ec2:Owner": "amazon"
}
}
},
{
"Effect": "Allow",
"Action": "ec2:RunInstances",
"Resource": [
"arn:aws:ec2:region:account:instance/*",
"arn:aws:ec2:region:account:subnet/*",
"arn:aws:ec2:region:account:volume/*",
"arn:aws:ec2:region:account:network-interface/*",
"arn:aws:ec2:region:account:key-pair/*",
"arn:aws:ec2:region:account:security-group/*"
]
}
]
}
```
b.Instance type

The following policy allows users to launch instances using only the t2.micro or t2.small instance type, which you might do to control costs. The users can't launch larger instances because the Condition element of the first statement tests whether ec2:InstanceType is either t2.micro or t2.small.
```
{
"Version": "2012-10-17",
"Statement": [{
"Effect": "Allow",
"Action": "ec2:RunInstances",
"Resource": [
"arn:aws:ec2:region:account:instance/*"
],
"Condition": {
"StringEquals": {
"ec2:InstanceType": ["t2.micro", "t2.small"]
}
}
},
{
"Effect": "Allow",
"Action": "ec2:RunInstances",
"Resource": [
"arn:aws:ec2:region::image/ami-*",
"arn:aws:ec2:region:account:subnet/*",
"arn:aws:ec2:region:account:network-interface/*",
"arn:aws:ec2:region:account:volume/*",
arn:aws:ec2:region:account:key-pair/*",
arn:aws:ec2:region:account:security-group/*"
]
}
]
}
```
Alternatively, you can create a policy that denies users permission to launch any instances except t2.micro and t2.small instance types.
```
{
```

```
"Version": "2012-10-17",
"Statement": [{
"Effect": "Deny",
"Action": "ec2:RunInstances",
"Resource": [
"arn:aws:ec2:region:account:instance/*"
],
"Condition": {
"StringNotEquals": {
"ec2:InstanceType": ["t2.micro", "t2.small"]
}
}
},
{
"Effect": "Allow",
"Action": "ec2:RunInstances",
"Resource": [
"arn:aws:ec2:region::image/ami-*",
"arn:aws:ec2:region:account:network-interface/*",
"arn:aws:ec2:region:account:instance/*",
"arn:aws:ec2:region:account:subnet/*",
"arn:aws:ec2:region:account:volume/*",
"arn:aws:ec2:region:account:key-pair/*",
arn:aws:ec2:region:account:security-group/*"
]
}
]
}
```

c.Subnet

The following policy allows users to launch instances using only the specified subnet, subnet-12345678. The group can't launch instances into any another subnet (unless another statement grants the users permission to do so). Users are still able to launch instances into EC2-Classic.

```
{
"Version": "2012-10-17",
"Statement": [{
"Effect": "Allow",
"Action": "ec2:RunInstances",
"Resource": [
"arn:aws:ec2:region:account:subnet/subnet-12345678",
"arn:aws:ec2:region:account:network-interface/*",
"arn:aws:ec2:region:account:instance/*",
"arn:aws:ec2:region:account:volume/*",
"arn:aws:ec2:region::image/ami-*",
"arn:aws:ec2:region:account:key-pair/*",
"arn:aws:ec2:region:account:security-group/*"
]
}
]
}
```

Alternatively, you could create a policy that denies users permission to launch an instance into any other subnet. The statement does this by denying permission to create a network interface, except where subnet subnet-12345678 is specified. This denial overrides any other policies that are created to allow launching instances into other subnets. Users are still able to launch instances into EC2-Classic.

```
{
"Version": "2012-10-17",
"Statement": [{
"Effect": "Deny",
"Action": "ec2:RunInstances",
"Resource": [
"arn:aws:ec2:region:account:network-interface/*"
],
"Condition": {
"ArnNotEquals": {
"ec2:Subnet": "arn:aws:ec2:region:account:subnet/subnet-12345678" }
}
},
```

```
{
"Effect": "Allow",
"Action": "ec2:RunInstances",
"Resource": [
"arn:aws:ec2:region::image/ami-*",
"arn:aws:ec2:region:account:network-interface/*",
"arn:aws:ec2:region:account:instance/*",
"arn:aws:ec2:region:account:subnet/*",
"arn:aws:ec2:region:account:volume/*",
"arn:aws:ec2:region:account:key-pair/*",
"arn:aws:ec2:region:account:security-group/*"
]
}
]
}
```

QUESTION 86
A 3-tier e-commerce web application is current deployed on-premises and will be migrated to AWS for greater scalability and elasticity The web server currently shares read-only data using a network distributed file system The app server tier uses a clustering mechanism for discovery and shared session state that depends on IP multicast The database tier uses shared-storage clustering to provide database fall over capability, and uses several read slaves for scaling Data on all servers and the distributed file system directory is backed up weekly to off-site tapes.
Which AWS storage and database architecture meets the requirements of the application?

A. Web servers: store read-only data in S3, and copy from S3 to root volume at boot time. App servers: share state using a combination of DynamoDB and IP unicast. Database: use RDS with multi-AZ deployment and one or more read replicas. Backup: web servers, app servers, and database backed up weekly to Glacier using snapshots.
B. Web servers: store read-only data in an EC2 NFS server; mount to each web server at boot time. App servers: share state using a combination of DynamoDB and IP multicast. Database: use RDS with multi-AZ deployment and one or more Read Replicas. Backup: web and app servers backed up weekly via AMIs, database backed up via DB snapshots.
C. Web servers: store read-only data in S3, and copy from S3 to root volume at boot time. App servers: share state using a combination of DynamoDB and IP unicast. Database: use RDS with multi-AZ deployment and one or more Read Replicas. Backup: web and app servers backed up weekly via AMIs, database backed up via DB snapshots.
D. Web servers: store read-only data in S3, and copy from S3 to root volume at boot time. App servers: share state using a combination of DynamoDB and IP unicast. Database: use RDS with multi-AZ deployment. Backup: web and app servers backed up weekly via AMIs, database backed up via DB snapshots.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Amazon RDS Multi-AZ deployments provide enhanced availability and durability for Database (DB) Instances, making them a natural fit for production database workloads. When you provision a Multi-AZ DB Instance, Amazon RDS automatically creates a primary DB Instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). Each AZ runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. In case of an infrastructure failure (for example, instance hardware failure, storage failure, or network disruption), Amazon RDS performs an automatic failover to the standby, so that you can resume database operations as soon as the failover is complete. Since the endpoint for your DB Instance remains the same after a failover, your application can resume database operation without the need for manual administrative intervention.
Benefits
Enhanced Durability
Multi-AZ deployments for the MySQL, Oracle, and PostgreSQL engines utilize synchronous physical replication to keep data on the standby up-to-date with the primary. Multi-AZ deployments for the SQL Server engine use synchronous logical replication to achieve the same result, employing SQL Server-native Mirroring technology. Both approaches safeguard your data in the event of a DB Instance failure or loss of an Availability Zone.
If a storage volume on your primary fails in a Multi-AZ deployment, Amazon RDS automatically initiates a failover to the up-to-date standby. Compare this to a Single-AZ deployment: in case of a Single-AZ database failure, a user-initiated point-in-time-restore operation will be required. This operation can take several hours to complete, and any data updates that occurred after the latest restorable time (typically within the last five minutes) will not be available.
Amazon Aurora employs a highly durable, SSD-backed virtualized storage layer purpose-built for database workloads.
Amazon Aurora automatically replicates your volume six ways, across three Availability Zones. Amazon Aurora storage is fault-tolerant, transparently handling the loss of up to two copies of data without affecting database write availability and up to three copies without affecting read availability. Amazon Aurora storage is also self-healing. Data blocks and disks are continuously scanned for errors and replaced automatically.
Increased Availability
You also benefit from enhanced database availability when running Multi-AZ deployments. If an Availability Zone failure or DB Instance failure occurs, your availability impact is limited to the time automatic failover takes to complete: typically under one minute for Amazon Aurora and one to two minutes for other database engines (see the RDS FAQ for details).
The availability benefits of Multi-AZ deployments also extend to planned maintenance and backups. In the case of system upgrades like OS patching or DB Instance scaling, these operations are applied first on the standby, prior to the automatic failover. As a result, your availability impact is, again, only the time required for automatic failover to

complete.
Unlike Single-AZ deployments, I/O activity is not suspended on your primary during backup for Multi-AZ deployments for the MySQL, Oracle, and PostgreSQL engines, because the backup is taken from the standby. However, note that you may still experience elevated latencies for a few minutes during backups for Multi-AZ deployments.
On instance failure in Amazon Aurora deployments, Amazon RDS uses RDS Multi-AZ technology to automate failover to one of up to 15 Amazon Aurora Replicas you have created in any of three Availability Zones. If no Amazon Aurora Replicas have been provisioned, in the case of a failure, Amazon RDS will attempt to create a new Amazon Aurora DB instance for you automatically.
No Administrative Intervention
DB Instance failover is fully automatic and requires no administrative intervention. Amazon RDS monitors the health of your primary and standbys, and initiates a failover automatically in response to a variety of failure conditions.
Failover conditions
Amazon RDS detects and automatically recovers from the most common failure scenarios for Multi-AZ deployments so that you can resume database operations as quickly as possible without administrative intervention. Amazon RDS automatically performs a failover in the event of any of the following:
Loss of availability in primary Availability Zone Loss of network connectivity to primary Compute unit failure on primary Storage failure on primary
Note: When operations such as DB Instance scaling or system upgrades like OS patching are initiated for Multi-AZ deployments, for enhanced availability, they are applied first on the standby prior to an automatic failover. As a result, your availability impact is limited only to the time required for automatic failover to complete. Note that Amazon RDS Multi-AZ deployments do not failover automatically in response to database operations such as long running queries, deadlocks or database corruption errors.

QUESTION 87
Your company plans to host a large donation website on Amazon Web Services (AWS). You anticipate a large and undetermined amount of traffic that will create many database writes. To be certain that you do not drop any writes to a database hosted on AWS.
Which service should you use?

A. Amazon RDS with provisioned IOPS up to the anticipated peak write throughput.
B. Amazon Simple Queue Service (SQS) for capturing the writes and draining the queue to write to the database.
C. Amazon ElastiCache to store the writes until the writes are committed to the database.
D. Amazon DynamoDB with provisioned write throughput up to the anticipated peak write throughput.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Amazon Simple Queue Service (Amazon SQS) offers a reliable, highly scalable hosted queue for storing messages as they travel between computers. By using Amazon SQS, developers can simply move data between distributed application components performing different tasks, without losing messages or requiring each component to be always available. Amazon SQS makes it easy to build a distributed, decoupled application, working in close conjunction with the Amazon Elastic Compute Cloud (Amazon EC2) and the other AWS infrastructure web services.
What can I do with Amazon SQS?
Amazon SQS is a web service that gives you access to a message queue that can be used to store messages while waiting for a computer to process them. This allows you to quickly build message queuing applications that can be run on any computer on the internet. Since Amazon SQS is highly scalable and you only pay for what you use, you can start small and grow your application as you wish, with no compromise on performance or reliability. This lets you focus on building sophisticated message-based applications, without worrying about how the messages are stored and managed. You can use Amazon SQS with software applications in various ways. For example, you can:
Integrate Amazon SQS with other AWS infrastructure web services to make applications more reliable and flexible.
Use Amazon SQS to create a queue of work where each message is a task that needs to be completed by a process. One or many computers can read tasks from the queue and perform them.
Build a microservices architecture, using queues to connect your microservices.
Keep notifications of significant events in a business process in an Amazon SQS queue. Each event can have a corresponding message in a queue, and applications that need to be aware of the event can read and process the messages.

QUESTION 88
You need a persistent and durable storage to trace call activity of an IVR (Interactive Voice Response) system. Call duration is mostly in the 2-3 minutes timeframe. Each traced call can be either active or terminated. An external application needs to know each minute the list of currently active calls. Usually there are a few calls/second, but once per month there is a periodic peak up to 1000 calls/second for a few hours. The system is open 24/7 and any downtime should be avoided. Historical data is periodically archived to files. Cost saving is a priority for this project.
What database implementation would better fit this scenario, keeping costs as low as possible?

A. Use DynamoDB with a "Calls" table and a Global Secondary Index on a "State" attribute that can equal to "active" or "terminated". In this way the Global Secondary Index can be used for all items in the table.
B. Use RDS Multi-AZ with a "CALLS" table and an indexed "STATE" field that can be equal to "ACTIVE" or 'TERMINATED". In this way the SQL query is optimized by the use of the Index.
C. Use RDS Multi-AZ with two tables, one for "ACTIVE.CALLS" and one for "TERMINATED.CALLS". In this way the "ACTIVE.CALLS" table is always small and effective to access.
D. Use DynamoDB with a "Calls" table and a Global Secondary Index on a "IsActive" attribute that is present for active calls only. In this way the Global Secondary Index is sparse and more effective.

Correct Answer: C
Section: (none)

Explanation

Explanation/Reference:

QUESTION 89
Your company hosts a social media site supporting users in multiple countries. You have been asked to provide a highly available design tor the application that leverages multiple regions tor the most recently accessed content and latency sensitive portions of the wet) site The most latency sensitive component of the application involves reading user preferences to support web site personalization and ad selection.
In addition to running your application in multiple regions, which option will support this application's requirements?

A. Serve user content from S3. CloudFront and use Route53 latency-based routing between ELBs in each region Retrieve user preferences from a local DynamoDB table in each region and leverage SQS to capture changes to user preferences with SOS workers for propagating updates to each table.
B. Use the S3 Copy API to copy recently accessed content to multiple regions and serve user content from S3. CloudFront with dynamic content and an ELB in each region Retrieve user preferences from an ElasticCache cluster in each region and leverage SNS notifications to propagate user preference changes to a worker node in each region.
C. Use the S3 Copy API to copy recently accessed content to multiple regions and serve user content from S3 CloudFront and Route53 latency-based routing Between ELBs In each region Retrieve user preferences from a DynamoDB table and leverage SQS to capture changes to user preferences with SOS workers for propagating DynamoDB updates.
D. Serve user content from S3. CloudFront with dynamic content, and an ELB in each region Retrieve user preferences from an ElastiCache cluster in each region and leverage Simple Workflow (SWF) to manage the propagation of user preferences from a centralized OB to each Elasticache cluster.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 90
You've been brought in as solutions architect to assist an enterprise customer with their migration of an e-commerce platform to Amazon Virtual Private Cloud (VPC) The previous architect has already deployed a 3-tier VPC.
The configuration is as follows:

```
VPC: vpc-2f8bc447
IGW: igw-2d8bc445
NACL: ad-208bc448
Subnets and Route Tables:
Web servers: subnet-258bc44d
Application servers: subnet-248bc44c
Database servers: subnet-9189c6f9
Route Tables:
rrb-218bc449
rtb-238bc44b
Associations:
subnet-258bc44d: rtb-218bc449
subnet-248bc44c : rtb-238bc44b
subnet-9189c6f9 : rtb-238bc44b
```

You are now ready to begin deploying EC2 instances into the VPC Web servers must have direct access to the internet Application and database servers cannot have direct access to the internet.
Which configuration below will allow you the ability to remotely administer your application and database servers, as well as allow these servers to retrieve updates from the Internet?

A. Create a bastion and NAT instance in subnet-258bc44d, and add a route from rtb- 238bc44b to the NAT instance.
B. Add a route from rtb-238bc44b to igw-2d8bc445 and add a bastion and NAT instance within subnet-248bc44c.
C. Create a bastion and NAT instance in subnet-248bc44c, and add a route from rtb- 238bc44b to subnet-258bc44d.
D. Create a bastion and NAT instance in subnet-258bc44d, add a route from rtb-238bc44b to Igw-2d8bc445, and a new NACL that allows access between subnet-258bc44d and subnet-248bc44c.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 91
You are designing a multi-platform web application for AWS The application will run on EC2 instances and will be accessed from PCs. Tablets and smart phones Supported accessing platforms are Windows, MacOS, IOS and Android Separate sticky session and SSL certificate setups are required for different platform types.
Which of the following describes the most cost effective and performance efficient architecture setup?

A. Setup a hybrid architecture to handle session state and SSL certificates on-prem and separate EC2 Instance groups running web applications for different platform types running in a VPC.
B. Set up one ELB for all platforms to distribute load among multiple instance under it Each EC2 instance implements ail functionality for a particular platform.
C. Set up two ELBs The first ELB handles SSL certificates for all platforms and the second ELB handles session stickiness for all platforms for each ELB run separate EC2 instance groups to handle the web application for each platform.
D. Assign multiple ELBS to an EC2 instance or group of EC2 instances running the common components of the web application, one ELB for each platform type Session stickiness and SSL termination are done at the ELBs.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 92
An administrator is using Amazon CloudFormation to deploy a three tier web application that consists of a web tier and application tier that will utilize Amazon DynamoDB for storage when creating the CloudFormation template.
Which of the following would allow the application instance access to the DynamoDB tables without exposing API credentials?

A. Create an Identity and Access Management Role that has the required permissions to read and write from the required DynamoDB table and associate the Role to the application instances by referencing an instance profile.
B. Use the Parameter section in the Cloud Formation template to nave the user input Access and Secret Keys from an already created IAM user that has me permissions required to read and write from the required DynamoDB table.
C. Create an Identity and Access Management Role that has the required permissions to read and write from the required DynamoDB table and reference the Role in the instance profile property of the application instance.
D. Create an identity and Access Management user in the CloudFormation template that has permissions to read and write from the required DynamoDB table, use the GetAtt function to retrieve the Access and secret keys and pass them to the application instance through user-data.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 93
Your company has recently extended its datacenter into a VPC on AWS to add burst computing capacity as needed Members of your Network Operations Center need to be able to go to the AWS Management Console and administer Amazon EC2 instances as necessary. You don't want to create new IAM users for each NOC member and make those users sign in again to the AWS Management Console.
Which option below will meet the needs for your NOC members?

A. Use OAuth 2.0 to retrieve temporary AWS security credentials to enable your NOC members to sign in to the AWS Management Console.
B. Use web Identity Federation to retrieve AWS temporary security credentials to enable your NOC members to sign in to the AWS Management Console.
C. Use your on-premises SAML 2.0-compliant identity provider (I DP) to grant the NOC members federated access to the AWS Management Console via the AWS single sign-on (SSO) endpoint.
D. Use your on-premises SAML2.0-compliam identity provider (IDP) to retrieve temporary security credentials to enable NOC members to sign in to the AWS Management Console.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 94
You have an application running on an EC2 Instance which will allow users to download flies from a private S3 bucket using a pre-signed URL. Before generating the URL the application should verify the existence of the file in S3.
How should the application use AWS credentials to access the S3 bucket securely?

A. Use the AWS account access Keys the application retrieves the credentials from the source code of the application.
B. Create an IAM user for the application with permissions that allow list access to the S3 bucket launch the instance as the I AM user and retrieve the IAM user's credentials from the EC2 instance user data.
C. Create an IAM role for EC2 that allows list access to objects in the S3 bucket. Launch the instance with the role, and retrieve the role's credentials from the EC2 Instance metadata
D. Create an IAM user for the application with permissions that allow list access to the S3 bucket. The application retrieves the IAM user credentials from a temporary directory with permissions that allow read access only to the application user.

Correct Answer: C
Section: (none)

Explanation

Explanation/Reference:

QUESTION 95
A benefits enrollment company is hosting a 3-tier web application running in a VPC on AWS which includes a NAT (Network Address Translation) instance in the public Web tier. There is enough provisioned capacity for the expected workload tor the new fiscal year benefit enrollment period plus some extra overhead Enrollment proceeds nicely for two days and then the web tier becomes unresponsive, upon investigation using Cloudwatch and other monitoring tools it is discovered that there is an extremely large and unanticipated amount of inbound traffic coming from a set of 15 specific IP addresses over port 80 from a country where the benefits company has no customers. The web tier instances are so overloaded that benefit enrollment administrators cannot even SSH into them.
Which activity would be useful in defending against this attack?

A. Create a custom route table associated with the web tier and block the attacking IP addresses from the IGW (Internet Gateway)
B. Change the EIP (Elastic IP Address) of the NAT instance in the web tier subnet and update the Main Route Table with the new EIP
C. Create 15 Security Group rules to block the attacking IP addresses over port 80
D. Create an inbound NACL (Network Access control list) associated with the web tier subnet with deny rules to block the attacking IP addresses

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Use AWS Identity and Access Management (IAM) to control who in your organization has permission to create and manage security groups and network ACLs (NACL). Isolate the responsibilities and roles for better defense. For example, you can give only your network administrators or security admin the permission to manage the security groups and restrict other roles.

QUESTION 96
You are developing a new mobile application and are considering storing user preferences in AWS.2w This would provide a more uniform cross-device experience to users using multiple mobile devices to access the application. The preference data for each user is estimated to be 50KB in size Additionally 5 million customers are expected to use the application on a regular basis.
The solution needs to be cost-effective, highly available, scalable and secure, how would you design a solution to meet the above requirements?

A. Setup an RDS MySQL instance in 2 availability zones to store the user preference data. Deploy a public facing application on a server in front of the database to manage security and access credentials
B. Setup a DynamoDB table with an item for each user having the necessary attributes to hold the user preferences. The mobile application will query the user preferences directly from the DynamoDB table. Utilize STS. Web Identity Federation, and DynamoDB Fine Grained Access Control to authenticate and authorize access.
C. Setup an RDS MySQL instance with multiple read replicas in 2 availability zones to store the user preference data .The mobile application will query the user preferences from the read replicas. Leverage the MySQL user management and access privilege system to manage security and access credentials.
D. Store the user preference data in S3 Setup a DynamoDB table with an item for each user and an item attribute pointing to the user' S3 object. The mobile application will retrieve the S3 URL from DynamoDB and then access the S3 object directly utilize STS, Web identity Federation, and S3 ACLs to authenticate and authorize access.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 97
You deployed your company website using Elastic Beanstalk and you enabled log file rotation to S3. An Elastic Map Reduce job is periodically analyzing the logs on S3 to build a usage dashboard that you share with your CIO.
You recently improved overall performance of the website using Cloud Front for dynamic content delivery and your website as the origin.
After this architectural change, the usage dashboard shows that the traffic on your website dropped by an order of magnitude.
How do you fix your usage dashboard'?

A. Enable Cloud Front to deliver access logs to S3 and use them as input of the Elastic Map Reduce job.
B. Turn on Cloud Trail and use trail log tiles on S3 as input of the Elastic Map Reduce job
C. Change your log collection process to use Cloud Watch ELB metrics as input of the Elastic Map Reduce job
D. Use Elastic Beanstalk "Rebuild Environment" option to update log delivery to the Elastic Map Reduce job.
E. Use Elastic Beanstalk Restart App server(s)" option to update log delivery to the Elastic Map Reduce job

Correct Answer: D
Section: (none)

Explanation

Explanation/Reference:

QUESTION 98
A web-startup runs its very successful social news application on Amazon EC2 with an Elastic Load Balancer, an Auto-Scaling group of Java/Tomcat application-servers, and DynamoDB as data store. The main web-application best runs on m2 x large instances since it is highly memory- bound Each new deployment requires semi-automated creation and testing of a new AMI for the application servers which takes quite a while ana is therefore only done once per week.
Recently, a new chat feature has been implemented in nodejs and wails to be integrated in the architecture. First tests show that the new component is CPU bound Because the company has some experience with using Chef, they decided to streamline the deployment process and use AWS Ops Works as an application lifecycle tool to simplify management of the application and reduce the deployment cycles.
What configuration in AWS Ops Works is necessary to integrate the new chat module in the most cost-efficient and flexible way?

A. Create one AWS OpsWorks stack, create one AWS Ops Works layer, create one custom recipe
B. Create one AWS OpsWorks stack create two AWS Ops Works layers, create one custom recipe
C. Create two AWS OpsWorks stacks create two AWS Ops Works layers, create one custom recipe
D. Create two AWS OpsWorks stacks create two AWS Ops Works layers, create two custom recipe

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 99
Select the correct set of options. These are the initial settings for the default security group:

A. Allow no inbound traffic, Allow all outbound traffic and Allow instances associated with this security group to talk to each other
B. Allow all inbound traffic, Allow no outbound traffic and Allow instances associated with this security group to talk to each other
C. Allow no inbound traffic, Allow all outbound traffic and Does NOT allow instances associated with this security group to talk to each other
D. Allow all inbound traffic. Allow all outbound traffic and Does NOT allow instances associated with this security group to talk to each other

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 100
How can an EBS volume that is currently attached to an EC2 instance be migrated from one Availability Zone to another?

A. Detach the volume and attach it to another EC2 instance in the other AZ.
B. Simply create a new volume in the other AZ and specify the original volume as the source.
C. Create a snapshot of the volume, and create a new volume from the snapshot in the other AZ.
D. Detach the volume, then use the ec2-migrate-volume command to move it to another AZ.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 101
After launching an instance that you intend to serve as a NAT (Network Address Translation) device in a public subnet you modify your route tables to have the NAT device be the target of internet bound traffic of your private subnet. When you try and make an outbound connection to the internet from an instance in the private subnet, you are not successful.
Which of the following steps could resolve the issue?

A. Disabling the Source/Destination Check attribute on the NAT instance

B. Attaching an Elastic IP address to the instance in the private subnet
C. Attaching a second Elastic Network Interface (ENI) to the NAT instance, and placing it in the private subnet
D. Attaching a second Elastic Network Interface (ENI) to the instance in the private subnet, and placing it in the public subnet

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
htto://docs.aws.amazon,com/worksoaces/latest/adminguide/gsg create vpc.html

QUESTION 102
Your application provides data transformation services. Files containing data to be transformed are first uploaded to Amazon S3 and then transformed by a fleet of spot EC2 instances. Files submitted by your premium customers must be transformed with the highest priority.
How should you implement such a system?

A. Use a DynamoDB table with an attribute defining the priority level. Transformation instances will scan the table for tasks, sorting the results by priority level.
B. Use Route 53 latency based-routing to send high priority tasks to the closest transformation instances.
C. Use two SQS queues, one for high priority messages, the other for default priority. Transformation instances first poll the high priority queue; if there is no message, they poll the default priority queue.
D. Use a single SQS queue. Each message contains the priority level. Transformation instances poll high-priority messages first.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 103
Which of the following are characteristics of Amazon VPC subnets? (Choose 2)

A. Each subnet spans at least 2 Availability Zones to provide a high-availability environment.
B. Each subnet maps to a single Availability Zone.
C. CIDR block mask of/25 is the smallest range supported.
D. By default, all subnets can route between each other, whether they are private or public.
E. Instances in a private subnet can communicate with the Internet only if they have an Elastic IP.

Correct Answer: AE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 104
In AWS, which security aspects are the customer's responsibility? (Choose 4)

A. Security Group and ACL (Access Control List) settings
B. Decommissioning storage devices
C. Patch management on the EC2 instance's operating system
D. Life-cycle management of IAM credentials
E. Controlling physical access to compute resources
F. Encryption of EBS (Elastic Block Storage) volumes

Correct Answer: ACDE
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
http://media.amazonwebservices.com/AWS Security Best Practices.pdf

QUESTION 105
When you put objects in Amazon S3, what is the indication that an object was successfully stored?

A. A HTTP 200 result code and MD5 checksum, taken together, indicate that the operation was successful.
B. Amazon S3 is engineered for 99.999999999% durability. Therefore there is no need to confirm that data was inserted.
C. A success code is inserted into the S3 object metadata.
D. Each S3 account has a special bucket named _s3_logs. Success codes are written to this bucket with a timestamp and checksum

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference: