Amazon.Premium.AWS Certified Solutions Architect - Professional.by.VCEplus.869q

Number: Amazon VCEplus
Passing Score: 800
Time Limit: 120 min
File Version: 20.4

Certification: AWS Certified Solutions Architect - Professional
Certification Full Name: AWS Certified Solutions Architect - Professional
Certification Provider: Amazon
Website: www.vceplus.com - www.vceplus.co - www.vceplus.io
Free Exam: https://vceplus.io/exam-aws-certified-solutions-architect-professional/
Questions & Answers Exam Engine is rigorously checked before being put up for sale. We make sure there is nothing irrelevant in AWS Certified Solutions Architect - Professional exam products and you get latest questions. We strive to deliver the best AWS Certified Solutions Architect - Professional exam product for top grades in your first attempt.

Exam A

QUESTION 1
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, change the storage class of the S3 objects to Reduced Redundancy Storage.
C. 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: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 2
An e-commerce company is revamping its IT infrastructure and is planning to use AWS services. The company's CIO has asked a Solutions Architect to design a simple, highly available, and loosely coupled order processing application. The application is responsible for receiving and processing orders before storing them in an Amazon DynamoDB table. The application has a sporadic traffic pattern and should be able to scale during marketing campaigns to process the orders with minimal delays.
Which of the following is the MOST reliable approach to meet the requirements?

A. Receive the orders in an Amazon EC2-hosted database and use EC2 instances to process them.
B. Receive the orders in an Amazon SQS queue and trigger an AWS Lambda function to process them.
C. Receive the orders using the AWS Step Functions program and trigger an Amazon ECS container to process them.
D. Receive the orders in Amazon Kinesis Data Streams and use Amazon EC2 instances to process them.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 3
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: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 4
A solutions architect needs to advise a company on how to migrate its on-premises data processing application to the AWS Cloud. Currently, users upload input files through a web portal. The web server then stores the uploaded files on NAS and messages the processing server over a message queue. Each media file can take up to 1 hour to process. The company has determined that the number of media files awaiting processing is significantly higher during business hours, with the number of files rapidly declining after business hours.
What is the MOST cost-effective migration recommendation?

A. Create a queue using Amazon SQS. Configure the existing web server to publish to the new queue. When there are messages in the queue, invoke an AWS Lambda function to pull requests from the queue and process the files. Store the processed files in an Amazon S3 bucket.
B. Create a queue using Amazon MQ. Configure the existing web server to publish to the new queue. When there are messages in the queue, create a new Amazon EC2 instance to pull requests from the queue and process the files. Store the processed files in Amazon EFS. Shut down the EC2 instance after the task is complete.
C. Create a queue using Amazon MQ. Configure the existing web server to publish to the new queue. When there are messages in the queue, invoke an AWS Lambda function to pull requests from the queue and process the files. Store the processed files in Amazon EFS.
D. Create a queue using Amazon SQS. Configure the existing web server to publish to the new queue. Use Amazon EC2 instances in an EC2 Auto Seating group to pull requests from the queue and process the files. Scale the EC2 instances based on the SQS queue length. Store the processed files in an Amazon S3 bucket.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 5
A company is moving a business-critical, multi-tier application to AWS. The architecture consists of a desktop client application and server infrastructure. The server infrastructure resides in an on-premises data center that frequently fails to maintain the application uptime SLA of 99.95%. A Solutions Architect must re-architect the application to ensure that it can meet or exceed the SLA.
The application contains a PostgreSQL database running on a single virtual machine. The business logic and presentation layers are load balanced between multiple virtual machines. Remote users complain about slow load times while using this latency-sensitive application.
Which of the following will meet the availability requirements with little change to the application while improving user experience and minimizing costs?

A. Migrate the database to a PostgreSQL database in Amazon EC2. Host the application and presentation layers in automatically scaled Amazon ECS containers behind an Application Load Balancer. Allocate an Amazon WorkSpaces WorkSpace for each end user to improve the user experience.
B. Migrate the database to an Amazon RDS Aurora PostgreSQL configuration. Host the application and presentation layers in an Auto Scaling configuration on Amazon EC2 instances behind an Application Load Balancer. Use Amazon AppStream 2.0 to improve the user experience.
C. Migrate the database to an Amazon RDS PostgreSQL Multi-AZ configuration. Host the application and presentation layers in automatically scaled AWS Fargate containers behind a Network Load Balancer. Use Amazon ElastiCache to improve the user experience.
D. Migrate the database to an Amazon Redshift cluster with at least two nodes. Combine and host the application and presentation layers in automatically scaled Amazon ECS containers behind an Application Load Balancer. Use Amazon CloudFront to improve the user experience.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 6
A company is creating a REST API to share information with six of its partners based in the United States. The company has created an Amazon API Gateway Regional endpoint. Each of the six partners will access the API once per day to post daily sales figures.
After initial deployment, the company observes 1,000 requests per second originating from 500 different IP addresses around the world. The company believes this traffic is originating from a botnet and wants to secure its API while minimizing cost.
Which approach should the company take to secure its API?

A. Create an Amazon CloudFront distribution with the API as the origin. Create an AWS WAF web ACL with a rule to block clients that submit more than five

requests per day. Associate the web ACL with the CloudFront distribution. Configure CloudFront with an origin access identity (OAI) and associate it with the distribution. Configure API Gateway to ensure only the OAI can run the POST method.

B. Create an Amazon CloudFront distribution with the API as the origin. Create an AWS WAF web ACL with a rule to block clients that submit more than five requests per day. Associate the web ACL with the CloudFront distribution. Add a custom header to the CloudFront distribution populated with an API key. Configure the API to require an API key on the POST method.

C. Create an AWS WAF web ACL with a rule to allow access to the IP addresses used by the six partners. Associate the web ACL with the API. Create a resource policy with a request limit and associate it with the API. Configure the API to require an API key on the POST method.

D. Create an AWS WAF web ACL with a rule to allow access to the IP addresses used by the six partners. Associate the web ACL with the API. Create a usage plan with a request limit and associate it with the API. Create an API key and add it to the usage plan.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 7
A Solutions Architect is responsible for redesigning a legacy Java application to improve its availability, data durability, and scalability. Currently, the application runs on a single high-memory Amazon EC2 instance. It accepts HTTP requests from upstream clients, adds them to an in-memory queue, and responds with a 200 status. A separate application thread reads items from the queue, processes them, and persists the results to an Amazon RDS MySQL instance. The processing time for each item takes 90 seconds on average, most of which is spent waiting on external service calls, but the application is written to process multiple items in parallel.
Traffic to this service is unpredictable. During periods of high load, items may sit in the internal queue for over an hour while the application processes the backlog. In addition, the current system has issues with availability and data loss if the single application node fails.
Clients that access this service cannot be modified. They expect to receive a response to each HTTP request they sendwithin 10 seconds before they will time out and retry the request.
Which approach would improve the availability and durability of the system while decreasing the processing latency and minimizing costs?

A. Create an Amazon API Gateway REST API that uses Lambda proxy integration to pass requests to an AWS Lambda function. Migrate the core processing code to a Lambda function and write a wrapper class that provides a handler method that converts the proxy events to the internal application data model and invokes the processing module.

B. Create an Amazon API Gateway REST API that uses a service proxy to put items in an Amazon SQS queue. Extract the core processing code from the existing application and update it to pull items from Amazon SQS instead of an in-memory queue. Deploy the new processing application to smaller EC2 instances within an Auto Scaling group that scales dynamically based on the approximate number of messages in the Amazon SQS queue.

C. Modify the application to use Amazon DynamoDB instead of Amazon RDS. Configure Auto Scaling for the DynamoDB table. Deploy the application within an Auto Scaling group with a scaling policy based on CPU utilization. Back the inmemory queue with a memory-mapped file to an instance store volume and periodically write that file to Amazon S3.

D. Update the application to use a Redis task queue instead of the in-memory queue. Build a Docker container image for the application. Create an Amazon ECS task definition that includes the application container and a separate container to host Redis. Deploy the new task definition as an ECS service using AWS

Fargate, and enable Auto Scaling.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://aws.amazon.com/blogs/database/introducing-amazon-elasticsearch-service-as-a-target-in-aws-database-migrationservice/

QUESTION 8
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 9
A company runs a popular public-facing ecommerce website. Its user base is growing quickly from a local market to a national market. The website is hosted in an on-premises data center with web servers and a MySQL database. The company wants to migrate its workload to AWS. A solutions architect needs to create a solution to:
Improve security
Improve reliability
Improve availability
Reduce latency
Reduce maintenance
Which combination of steps should the solutions architect take to meet these requirements? (Choose three.)

A. Use Amazon EC2 instances in two Availability Zones for the web servers in an Auto Scaling group behind an Application Load Balancer.
B. Migrate the database to a Multi-AZ Amazon Aurora MySQL DB cluster.
C. Use Amazon EC2 instances in two Availability Zones to host a highly available MySQL database cluster.
D. Host static website content in Amazon S3. Use S3 Transfer Acceleration to reduce latency while serving webpages. Use AWS WAF to improve website security.
E. Host static website content in Amazon S3. Use Amazon CloudFront to reduce latency while serving webpages. Use AWS WAF to improve website security.
F. Migrate the database to a single-AZ Amazon RDS for MySQL DB instance.

Correct Answer: DEF
Section: (none)
Explanation

Explanation/Reference:

QUESTION 10
An auction website enables users to bid on collectible items. The auction rules require that each bid is processed only once and in the order it was received. The

current implementation is based on a fleet of Amazon EC2 web servers that write bid records into Amazon Kinesis Data Streams. A single t2.large instance has a cron job that runs the bid processor, which reads incoming bids from Kinesis Data Streams and processes each bid. The auction site is growing in popularity, but users are complaining that some bids are not registering.
Troubleshooting indicates that the bid processor is too slow during peak demand hours, sometimes crashes while processing, and occasionally loses track of which records is being processed. What changes should make the bid processing more reliable?

A. Refactor the web application to use the Amazon Kinesis Producer Library (KPL) when posting bids to Kinesis Data Streams. Refactor the bid processor to flag each record in Kinesis Data Streams as being unread, processing, and processed. At the start of each bid processing run, scan Kinesis Data Streams for unprocessed records.
B. Refactor the web application to post each incoming bid to an Amazon SNS topic in place of Kinesis Data Streams. Configure the SNS topic to trigger an AWS Lambda function that processes each bid as soon as a user submits it.
C. Refactor the web application to post each incoming bid to an Amazon SQS FIFO queue in place of Kinesis Data Streams. Refactor the bid processor to continuously the SQS queue. Place the bid processing EC2 instance in an Auto Scaling group with a minimum and a maximum size of 1.
D. Switch the EC2 instance type from t2.large to a larger general compute instance type. Put the bid processor EC2 instances in an Auto Scaling group that scales out the number of EC2 instances running the bid processor, based on the IncomingRecords metric in Kinesis Data Streams.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://d0.awsstatic.com/whitepapers/Building_a_Real_Time_Bidding_Platform_on_AWS_v1_Final.pdf

QUESTION 11
A company is running a web application with On-Demand Amazon EC2 instances in Auto Scaling groups that scale dynamically based on custom metrics. After extensive testing, the company determines that the m5.2xlarge instance size is optimal for the workload. Application data is stored in db.r4.4xlarge Amazon RDS instances that are confirmed to be optimal.
The traffic to the web application spikes randomly during the day.
What other cost-optimization methods should the company implement to further reduce costs without impacting the reliability of the application?

A. Double the instance count in the Auto Scaling groups and reduce the instance size to m5.large.
B. Reserve capacity for the RDS database and the minimum number of EC2 instances that are constantly running.
C. Reduce the RDS instance size to db.r4.xlarge and add five equivalently sized read replicas to provide reliability.
D. Reserve capacity for all EC2 instances and leverage Spot Instance pricing for the RDS database.

Correct Answer: B

Section: (none)
Explanation

Explanation/Reference:

QUESTION 12
A company runs a serverless application in a single AWS Region. The application accesses external URLs and extracts metadata from those sites. The company uses an Amazon Simple Notification Service (Amazon SNS) topic to publish URLs to an Amazon Simple Queue Service (Amazon SQS) queue. An AWS Lambda function uses the queue as an event source and processes the URLs from the queue. Results are saved to an Amazon S3 bucket.
The company wants to process each URL in other Regions to compare possible differences in site localization. URLs must be published from the existing Region. Results must be written to the existing S3 bucket in the current Region.
Which combination of changes will produce multi-Region deployment that meets these requirements? (Choose two.)

A. Deploy the SQS queue with the Lambda function to other Regions.
B. Subscribe the SNS topic in each Region to the SQS queue.
C. Subscribe the SQS queue in each Region to the SNS topic.
D. Configure the SQS queue to publish URLs to SNS topics in each Region.
E. Deploy the SNS topic and the Lambda function to other Regions.

Correct Answer: CE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 13
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 14
You have been given the task to define multiple AWS Data Pipeline schedules for different activities in the same pipeline.
Which of the following would successfully accomplish this task?

A. Creating multiple pipeline definition files
B. Defining multiple pipeline definitions in your schedule objects file and associating the desired schedule to the correct activity via its schedule field
C. Defining multiple schedule objects in your pipeline definition file and associating the desired schedule to the correct activity via its schedule field
D. Defining multiple schedule objects in the schedule field

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
To define multiple schedules for different activities in the same pipeline, in AWS Data Pipeline, you should define multiple schedule objects in your pipeline definition file and associate the desired schedule to the correct activity via its schedule field.
As an example of this, it could allow you to define a pipeline in which log files are stored in Amazon S3 each hour to drive generation of an aggregate report once a day.
Reference:
https://aws.amazon.com/datapipeline/faqs/

QUESTION 15

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 16
A company provides a centralized Amazon EC2 application hosted in a single shared VPC. The centralized application must be accessible from client applications running in the VPCs of other business units. The centralized application front end is configured with a Network Load Balancer (NLB) for scalability. Up to 10 business unit VPCs will need to be connected to the shared VPC. Some of the business unit VPC CIDR blocks overlap with the shared VPC, and some overlap with each other. Network connectivity to the centralized application in the shared VPC should be allowed from authorized business unit VPCs only.
Which network configuration should a solutions architect use to provide connectivity from the client applications in the business unit VPCs to the centralized application in the shared VPC?

A. Create an AWS Transit Gateway. Attach the shared VPC and the authorized business unit VPCs to the transit gateway.
Create a single transit gateway route table and associate it with all of the attached VPCs. Allow automatic propagation of routes from the attachments into the route table. Configure VPC routing tables to send traffic to the transit gateway
B. Create a VPC endpoint service using the centralized application NLB and enable the option to require endpoint acceptance. Create a VPC endpoint in each of the business unit VPCs using the service name of the endpoint service.
Accept authorized endpoint requests from the endpoint service console.
C. Create a VPC peering connection from each business unit VPC to the shared VPAccept the VPC peering connections from the shared VPC console.
Configure VPC routing tables to send traffic to the VPC peering connection.
D. Configure a virtual private gateway for the shared VPC and create customer gateways for each of the authorized business unit VPCs. Establish a Site-to-Site VPN connection from the business unit VPCs to the shared VPC. Configure VPC routing tables to send traffic to the VPN connection.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://d1.awsstatic.com/whitepapers/building-a-scalable-and-secure-multi-vpc-aws-network-infrastructure.pdf

QUESTION 17
One of the components that is part of ec2-net-utils used with ENI's is ec2ifscan.
Which of the following is not correct about ec2-net-utils?

A. ec2-net-utils generates an interface configuration file suitable for use with DHCP.
B. ec2-net-utils extends the functionality of the standard if up.
C. ec2-net-utils detaches a primary network interface from an instance.
D. ec2-net-utils identifies network interfaces when they are attached, detached, or reattached to a running instance.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Each instance in a VPC has a default elastic network interface (the primary network interface) that is assigned a private IP address from the IP address range of your VPC. You cannot detach a primary network interface from an instance. You can create and attach additional elastic network interfaces. Amazon Linux AMIs may contain additional scripts installed by AWS, known as ec2-net-utils. One of the components that is part of ec2-net-utils used with ENI's is ec2ifscan. Its function is to check for network interfaces that have not been configured and configure them.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html

QUESTION 18
Which AWS instance address has the following characteristics? :"If you stop an instance, its Elastic IP address is unmapped, and you must remap it when you restart the instance."

A. Both A and B
B. None of these
C. VPC Addresses
D. EC2 Addresses

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Stopping an instance EC2-Classic
If you stop an instance, its Elastic IP address is disassociated, and you must reassociate the Elastic IP address when you restart the instance.
EC2-VPC
If you stop an instance, its Elastic IP address remains associated.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-ip-addresses-eip.html

QUESTION 19
A 3-Ber e-commerce web application is currently deployed on-premises, and will be migrated to AWS for greater scalability and elasticity. The web tier 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 sharedstorage clustering to provide database failover 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: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Amazon Glacier doesn't suit all storage situations. Listed following are a few storage needs for which you should consider other AWS storage options instead of

Amazon Glacier.
Data that must be updated very frequently might be better served by a storage solution with lower read/write latencies, such as Amazon EBS, Amazon RDS, Amazon DynamoDB, or relational databases running on EC2.
Reference:
https://d0.awsstatic.com/whitepapers/Storage/AWS%20Storage%20Services%20Whitepaper-v9.pdf

QUESTION 20
A company recently completed a large-scale migration to AWS. Development teams that support various business units have their own accounts in AWS Organizations. A central cloud team is responsible for controlling which services and resources can be accessed, and for creating operational strategies for all teams within the company. Some teams are approaching their account service quotas. The cloud team needs to create an automated and operationally efficient solution to proactively monitor service quotas. Monitoring should occur every 15 minutes and send alerts when a team exceeds 80% utilization.
Which solution will meet these requirements?

A. Create a scheduled AWS Config rule to trigger an AWS Lambda function to call the GetServiceQuota API. If any service utilization is above 80%, publish a message to an Amazon Simple Notification Service (Amazon SNS) topic to alert the cloud team. Create an AWS CloudFormation template and deploy the necessary resources to each account.
B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that triggers an AWS Lambda function to refresh the AWS Trusted Advisor service limits checks and retrieve the most current utilization and service limit data. If the current utilization is above 80%, publish a message to an Amazon Simple Notification Service (Amazon SNS) topic to alert the cloud team. Create AWS CloudFormation StackSets that deploy the necessary resources to all Organizations accounts.
C. Create an Amazon CloudWatch alarm that triggers an AWS Lambda function to call the Amazon CloudWatch GetInsightRuleReport API to retrieve the most current utilization and service limit data. If the current utilization is above 80%, publish an Amazon Simple Email Service (Amazon SES) notification to alert the cloud team. Create AWS CloudFormation StackSets that deploy the necessary resources to all Organizations accounts.
D. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that triggers an AWS Lambda function to refresh the AWS Trusted Advisor service limits checks and retrieve the most current utilization and service limit data. If the current utilization is above 80%, use Amazon Pinpoint to send an alert to the cloud team. Create an AWS CloudFormation template and deploy the necessary resources to each account.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/solutions/implementations/limit-monitor/

QUESTION 21
A development team has created a new flight tracker application that provides near-real-time data to users. The application has a front end that consists of an Application Load Balancer (ALB) in front of two large Amazon EC2 instances in a single Availability Zone. Data is stored in a single Amazon RDS MySQL DB instance. An Amazon Route 53 DNS record points to the ALB.
Management wants the development team to improve the solution to achieve maximum reliability with the least amount of operational overhead.

Which set of actions should the team take?

A. Create RDS MySQL read replicas. Deploy the application to multiple AWS Regions. Use a Route 53 latency-based routing policy to route to the application.
B. Configure the DB instance as Multi-AZ. Deploy the application to two additional EC2 instances in different Availability Zones behind an ALB.
C. Replace the DB instance with Amazon DynamoDB global tables. Deploy the application in multiple AWS Regions. Use a Route 53 latency-based routing policy to route to the application.
D. Replace the DB instance with Amazon Aurora with Aurora Replicas. Deploy the application to multiple smaller EC2 instances across multiple Availability Zones in an Auto Scaling group behind an ALB.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZ.html

QUESTION 22
An enterprise company wants to allow its developers to purchase third-party software through AWS Marketplace. The company uses an AWS Organizations account structure with full features enabled, and has a shared services account in each organizational unit (OU) that will be used by procurement managers. The procurement team's policy indicates that developers should be able to obtain third-party software from an approved list only and use Private Marketplace in AWS Marketplace to achieve this requirement. The procurement team wants administration of Private Marketplace to be restricted to a role named procurement-manager-role, which could be assumed by procurement managers. Other IAM users, groups, roles, and account administrators in the company should be denied Private Marketplace administrative access.
What is the MOST efficient way to design an architecture to meet these requirements?

A. Create an IAM role named procurement-manager-role in all AWS accounts in the organization. Add the PowerUserAccess managed policy to the role. Apply an inline policy to all IAM users and roles in every AWS account to deny permissions on the AWSPrivateMarketplaceAdminFullAccess managed policy.
B. Create an IAM role named procurement-manager-role in all AWS accounts in the organization. Add the AdministratorAccess managed policy to the role. Define a permissions boundary with the AWSPrivateMarketplaceAdminFullAccess managed policy and attach it to all the developer roles.
C. Create an IAM role named procurement-manager-role in all the shared services accounts in the organization. Add the AWSPrivateMarketplaceAdminFullAccess managed policy to the role. Create an organization root-level SCP to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role. Create another organization root-level SCP to deny permissions to create an IAM role named procurement-manager-role to everyone in the organization.
D. Create an IAM role named procurement-manager-role in all AWS accounts that will be used by developers. Add the AWSPrivateMarketplaceAdminFullAccess managed policy to the role. Create an SCP in Organizations to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role. Apply the SCP to all the shared services accounts in the organization.

Correct Answer: D

Section: (none)
Explanation

Explanation/Reference:

QUESTION 23
True or false: In CloudFormation, you cannot create an Amazon RDS DB instance from a snapshot.

A. False, you can specify it in attributes
B. False, you can specify it in condition
C. False, you can specify it in resource properties
D. True

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In AWS CloudFormation, resource properties are additional options that you can specify on a resource. For example, you can specify the DB snapshot property for an Amazon RDS DB instance in order to create a DB instance from a snapshot.
Reference: http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/concept-resources.html

QUESTION 24
The MySecureData company has five branches across the globe. They want to expand their data centers such that their web server will be in the AWS and each branch would have their own database in the local data center. Based on the user login, the company wants to connect to the data center. How can MySecureData company implement this scenario with the AWS VPC?

A. Create five VPCs with the public subnet for the app server and setup the VPN gateway for each VPN to connect them individually.
B. Use the AWS VPN CloudHub to communicate with multiple VPN connections.
C. Use the AWS CloudGateway to communicate with multiple VPN connections.
D. It is not possible to connect different data centers from a single VPC.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. If the user wants to connect VPC from his own data centre, he can setup a public and VPN only subnet which uses hardware VPN access to connect with his data centre. If the organization has multiple VPN connections, he can provide secure communication between sites using the AWS VPN CloudHub.
The VPN CloudHub operates on a simple hub-and-spoke model that the user can use with or without a VPC. This design is suitable for customers with multiple branch offices and existing internet connections who would like to implement a convenient, potentially low-cost hub-and- spoke model for primary or backup connectivity between remote offices.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPN_CloudHub.html

QUESTION 25
In the context of AWS Cloud Hardware Security Module(HSM), does your application need to reside in the same VPC as the CloudHSM instance?

A. No, but the server or instance on which your application and the HSM client is running must have network (IP) reachability to the HSM.
B. Yes, always
C. No, but they must reside in the same Availability Zone.
D. No, but it should reside in same Availability Zone as the DB instance.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Your application does not need to reside in the same VPC as the CloudHSM instance. However, the server or instance on which your application and the HSM client is running must have network (IP) reachability to the HSM. You can establish network connectivity in a variety of ways, including operating your application in the same VPC, with VPC peering, with a VPN connection, or with Direct Connect.
Reference: https://aws.amazon.com/cloudhsm/faqs/

QUESTION 26
A company is moving a business-critical application onto AWS. It is a traditional three-tier web application using an Oracle database. Data must be encrypted in transit and at rest. The database hosts 12 TB of data. Network connectivity to the source Oracle database over the internal is allowed, and the company wants to reduce operational costs by using AWS Managed Services where possible. All resources within the web and application tiers have been migrated. The database has a few tables and a simple schema using primary keys only; however, it contains many Binary Large Object (BLOB) fields. It was not possible to use the database's native replication tools because of licensing restrictions.
Which database migration solution will result in the LEAST amount of impact to the application's availability?

A. Provision an Amazon RDS for Oracle instance. Host the RDS database within a virtual private cloud (VPC) subnet with internet access, and set up the RDS database as an encrypted Read Replica of the source database. Use SSL to encrypt the connection between the two databases. Monitor the replication performance by watching the RDS ReplicaLag metric.
During the application maintenance window, shut down the on-premises database and switch over the application connection to the RDS instance when there

is no more replication lag. Promote the Read Replica into a standalone database instance.

B. Provision an Amazon EC2 instance and install the same Oracle database software. Create a backup of the source database using the supported tools. During the application maintenance window, restore the backup into the Oracle database running in the EC2 instance. Set up an Amazon RDS for Oracle instance, and create an import job between the databases hosted in AWS. Shut down the source database and switch over the database connections to the RDS instance when the job is complete.

C. Use AWS DMS to load and replicate the dataset between the on-premises Oracle database and the replication instance hosted on AWS. Provision an Amazon RDS for Oracle instance with Transparent Data Encryption (TDE) enabled and configure it as a target for the replication instance. Create a customer-managed AWS KMS master key to set it as the encryption key for the replication instance. Use AWS DMS tasks to load the data into the target RDS instance. During the application maintenance window and after the load tasks reach the ongoing replication phase, switch the database connections to the new database.

D. Create a compressed full database backup of the on-premises Oracle database during an application maintenance window. While the backup is being performed, provision a 10 Gbps AWS Direct Connect connection to increase the transfer speed of the database backup files to Amazon S3, and shorten the maintenance window period. Use SSL/TLS to copy the files over the Direct Connect connection. When the backup files are successfully copied, start the maintenance window, and rise any of the Amazon RDS supported tools to import the data into a newly provisioned Amazon RDS for Oracle instance with encryption enabled. Wait until the data is fully loaded and switch over the database connections to the new database.
Delete the Direct Connect connection to cut unnecessary charges.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/blogs/apn/oracle-database-encryption-options-on-amazon-rds/
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Appendix.Oracle.Options.AdvSecurity.htm(DMS in transit encryption)
https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Security.html

QUESTION 27
A company is planning a large event where a promotional offer will be introduced. The company's website is hosted on AWS and backed by an Amazon RDS for PostgreSQL DB instance. The website explains the promotion and includes a sign-up page that collects user information and preferences. Management expects large and unpredictable volumes of traffic periodically, which will create many database writes. A solutions architect needs to build a solution that does not change the underlying data model and ensures that submissions are not dropped before they are committed to the database.
Which solution meets these requirements?

A. Immediately before the event, scale up the existing DB instance to meet the anticipated demand. Then scale down after the event.

B. Use Amazon SQS to decouple the application and database layers. Configure an AWS Lambda function to write items from the queue into the database.

C. Migrate to Amazon DynamoDB and manage throughput capacity with automatic scaling.

D. Use Amazon ElastiCache for Memcached to increase write capacity to the DB instance.

Correct Answer: D

Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/elasticache/faqs/

QUESTION 28
A company wants to migrate its data analytics environment from on premises to AWS. The environment consists of two simple Node.js applications. One of the applications collects sensor data and loads it into a MySQL database. The other application aggregates the data into reports. When the aggregation jobs run, some of the load jobs fail to run correctly.
The company must resolve the data loading issue. The company also needs the migration to occur without interruptions or changes for the company's customers.
What should a solutions architect do to meet these requirements?

A. Set up an Amazon Aurora MySQL database as a replication target for the on-premises database. Create an Aurora Replica for the Aurora MySQL database, and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as AWS Lambda functions behind a Network Load Balancer (NLB), and use Amazon RDS Proxy to write to the Aurora MySQL database. When the databases are synced, disable the replication job and restart the Aurora Replica as the primary instance. Point the collector DNS record to the NLB.
B. Set up an Amazon Aurora MySQL database. Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora. Move the aggregation jobs to run against the Aurora MySQL database. Set up collection endpoints behind an Application Load Balancer (ALB) as Amazon EC2 instances in an Auto Scaling group. When the databases are synced, point the collector DNS record to the ALDisable the AWS DMS sync task after the cutover from on premises to AWS.
C. Set up an Amazon Aurora MySQL database. Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora. Create an Aurora Replica for the Aurora MySQL database, and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as AWS Lambda functions behind an Application Load Balancer (ALB), and use Amazon RDS Proxy to write to the Aurora MySQL database. When the databases are synced, point the collector DNS record to the ALB. Disable the AWS DMS sync task after the cutover from on premises to AWS.
D. Set up an Amazon Aurora MySQL database. Create an Aurora Replica for the Aurora MySQL database, and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as an Amazon Kinesis data stream. Use Amazon Kinesis Data Firehose to replicate the data to the Aurora MySQL database. When the databases are synced, disable the replication job and restart the Aurora Replica as the primary instance. Point the collector DNS record to the Kinesis data stream.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 29
A company receives clickstream data files to Amazon S3 every five minutes. A Python script runs as a cron job once a day on an Amazon EC2 instance to process each file and load it into a database hosted on Amazon RDS. The cron job takes 15 to 30 minutes to process 24 hours of data. The data consumers ask for the data be available as soon as possible.
Which solution would accomplish the desired outcome?

A. Increase the size of the instance to speed up processing and update the schedule to run once an hour.
B. Convert the cron job to an AWS Lambda function and trigger this new function using a cron job on an EC2 instance.
C. Convert the cron job to an AWS Lambda function and schedule it to run once an hour using Amazon CloudWatch Events.
D. Create an AWS Lambda function that runs when a file is delivered to Amazon S3 using S3 event notifications.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/lambda/latest/dg/with-s3.html

QUESTION 30
A Solutions Architect is designing the storage layer for a recently purchased application. The application will be running on Amazon EC2 instances and has the following layers and requirements: Data layer: A POSIX file system shared across many systems.
Service layer: Static file content that requires block storage with more than 100k IOPS.
Which combination of AWS services will meet these needs? (Choose two.)

A. Data layer – Amazon S3
B. Data layer – Amazon EC2 Ephemeral Storage
C. Data layer – Amazon EFS
D. Service layer – Amazon EBS volumes with Provisioned IOPS
E. Service layer – Amazon EC2 Ephemeral Storage

Correct Answer: CE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 31

AnyCompany has acquired numerous companies over the past few years. The CIO for AnyCompany would like to keep the resources for each acquired company separate. The CIO also would like to enforce a chargeback model where each company pays for the AWS services it uses.
The Solutions Architect is tasked with designing an AWS architecture that allows AnyCompany to achieve the following:
Implementing a detailed chargeback mechanism to ensure that each company pays for the resources it uses.
AnyCompany can pay for AWS services for all its companies through a single invoice.
Developers in each acquired company have access to resources in their company only.
Developers in an acquired company should not be able to affect resources in their company only. A single identity store is used to authenticate Developers across all companies. Which of the following approaches would meet these requirements? (Choose two.)

A. Create a multi-account strategy with an account per company. Use consolidated billing to ensure that AnyCompany needs to pay a single bill only.
B. Create a multi-account strategy with a virtual private cloud (VPC) for each company. Reduce impact across companies by not creating any VPC peering links. As everything is in a single account, there will be a single invoice. Use tagging to create a detailed bill for each company.
C. Create IAM users for each Developer in the account to which they require access. Create policies that allow the users access to all resources in that account. Attach the policies to the IAM user.
D. Create a federated identity store against the company's Active Directory. Create IAM roles with appropriate permissions and set the trust relationships with AWS and the identity store. Use AWS STS to grant users access based on the groups they belong to in the identity store.
E. Create a multi-account strategy with an account per company. For billing purposes, use a tagging solution that uses a tag to identify the company that creates each resource.

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 32
What is the maximum length for a certificate ID in AWS IAM?

A. 1024 characters
B. 512 characters
C. 64 characters
D. 128 characters

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

Explanation:
The maximum length for a certificate ID is 128 characters.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/LimitationsOnEntities.html

QUESTION 33
A company has a web application that allows users to upload short videos. The videos are stored on Amazon EBS volumes and analyzed by custom recognition software for categorization.
The website contains static content that has variable traffic with peaks in certain months. The architecture consists of Amazon EC2 instances running in an Auto Scaling group for the web application and EC2 instances running in an Auto Scaling group to process an Amazon SQS-queue. The company wants to re-architect the application to reduce operational overhead using AWS managed services where possible and remove dependencies on third-party software.
Which solution meets these requirements?

A. Use Amazon ECS containers for the web application and Spot instances for the Scaling group that processes the SQS queue. Replace the custom software with Amazon Rekognition to categorize the videos.
B. Store the uploaded videos in Amazon EFS and mount the file system to the EC2 instances for the web application. Process the SQS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos.
C. Host the web application in Amazon S3. Store the uploaded videos in Amazon S3. Use S3 event notification to publish events to the SQS queue. Process the SQS queue with an AWS Lambda function that call the Amazon Rekognition API to categorize the videos.
D. Use AWS Elastic Beanstalk to launch EC2 instances in an Auto Scaling group for the application and launch a worker environment to process the SQS queue. Replace the custom software with Amazon Rekognition to categorize the videos.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 34
A company needs to store and process image data that will be uploaded from mobile devices using a custom mobile app.
Usage peaks between 8 AM and 5 PM on weekdays, with thousands of uploads per minute. The app is rarely used at any other time. A user is notified when image processing is complete.
Which combination of actions should a solutions architect take to ensure image processing can scale to handle the load?
(Choose three.)

A. Upload files from the mobile software directly to Amazon S3. Use S3 event notifications to create a message in an Amazon MQ queue.
B. Upload files from the mobile software directly to Amazon S3. Use S3 event notifications to create a message in an Amazon Simple Queue Service (Amazon SQS) standard queue.
C. Invoke an AWS Lambda function to perform image processing when a message is available in the queue.
D. Invoke an S3 Batch Operations job to perform image processing when a message is available in the queue.

E. Send a push notification to the mobile app by using Amazon Simple Notification Service (Amazon SNS) when processing is complete.
F. Send a push notification to the mobile app by using Amazon Simple Email Service (Amazon SES) when processing is complete.

Correct Answer: BEF
Section: (none)
Explanation

Explanation/Reference:

QUESTION 35
In DynamoDB, "The data is eventually consistent" means that_____.

A. a read request immediately after a write operation might not show the latest change.
B. a read request immediately after a write operation shows the latest change.
C. a write request immediately after a read operation might cause data loss.
D. a read request immediately after a write operation might cause data loss.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In DynamoDB, it takes time for the update to propagate to all copies. The data is eventually consistent, meaning that a read request immediately after a write operation might not show the latest change.
Reference: http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/APISummary.html

QUESTION 36
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: http://docs.aws.amazon.com/workspaces/latest/adminguide/gsg_create_vpc.html

QUESTION 37
An online e-commerce business is running a workload on AWS. The application architecture includes a web tier, an application tier for business logic, and a database tier for user and transactional data management. The database server has a 100 GB memory requirement. The business requires cost-efficient disaster recovery for the application with an RTO of 5 minutes and an RPO of 1 hour. The business also has a regulatory for out-of-region disaster recovery with a minimum distance between the primary and alternate sites of 250 miles.
Which of the following options can the Solutions Architect design to create a comprehensive solution for this customer that meets the disaster recovery requirements?

A. Back up the application and database data frequently and copy them to Amazon S3. Replicate the backups using S3 cross-region replication, and use AWS CloudFormation to instantiate infrastructure for disaster recovery and restore data from Amazon S3.
B. Employ a pilot light environment in which the primary database is configured with mirroring to build a standby database on m4.large in the alternate region. Use AWS CloudFormation to instantiate the web servers, application servers and load balancers in case of a disaster to bring the application up in the alternate region. Vertically resize the database to meet the full production demands, and use Amazon Route 53 to switch traffic to the alternate region.
C. Use a scaled-down version of the fully functional production environment in the alternate region that includes one instance of the web server, one instance of the application server, and a replicated instance of the database server in standby mode.
Place the web and the application tiers in an Auto Scaling behind a load balancer, which can automatically scale when the load arrives to the application. Use Amazon Route 53 to switch traffic to the alternate region.
D. Employ a multi-region solution with fully functional web, application, and database tiers in both regions with equivalent capacity. Activate the primary database in one region only and the standby database in the other region. Use Amazon Route 53 to automatically switch traffic from one region to another using health check routing policies.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 38
In Amazon IAM, what is the maximum length for a role name?

A. 128 characters

B. 512 characters
C. 64 characters
D. 256 characters

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In Amazon IAM, the maximum length for a role name is 64 characters.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/LimitationsOnEntities.html

QUESTION 39
A company is launching a new web application on Amazon EC2 instances. Development and production workloads exist in separate AWS accounts. According to the company's security requirements, only automated configuration tools are allowed to access the production account. The company's security team wants to receive immediate notification if any manual access to the production AWS account or EC2 instances occurs.
Which combination of actions should a solutions architect take in the production account to meet these requirements?
(Choose three.)

A. Turn on AWS CloudTrail logs in the application's primary AWS Region. Use Amazon Athena to query the logs for AwsConsoleSignIn events.
B. Configure Amazon Simple Email Service (Amazon SES) to send email to the security team when an alarm is activated.
C. Deploy EC2 instances in an Auto Scaling group. Configure the launch template to deploy instances without key pairs. Configure Amazon CloudWatch Logs to capture system access logs. Create an Amazon CloudWatch alarm that is based on the logs to detect when a user logs in to an EC2 instance.
D. Configure an Amazon Simple Notification Service (Amazon SNS) topic to send a message to the security team when an alarm is activated.
E. Turn on AWS CloudTrail logs for all AWS Regions. Configure Amazon CloudWatch alarms to provide an alert when an AwsConsoleSignIn event is detected.
F. Deploy EC2 instances in an Auto Scaling group. Configure the launch template to delete the key pair after launch. Configure Amazon CloudWatch Logs for the system access logs. Create an Amazon CloudWatch dashboard to show user logins over time.

Correct Answer: CDE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 40
A company is running an application on several Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer. The load on the

application varies throughout the day, and EC2 instances are scaled in and out on a regular basis.
Log files from the EC2 instances are copied to a central Amazon S3 bucket every 15 minutes. The security team discovers that log files are missing from some of the terminated EC2 instances.
Which set of actions will ensure that log files are copied to the central S3 bucket from the terminated EC2 instances?

A. Create a script to copy log files to Amazon S3, and store the script in a file on the EC2 instance. Create an Auto Scaling lifecycle hook and an Amazon EventBridge (Amazon CloudWatch Events) rule to detect lifecycle events from the Auto Scaling group. Invoke an AWS Lambda function on the autoscaling:EC2_INSTANCE_TERMINATING transition to send ABANDON to the Auto Scaling group to prevent termination run the script to copy the log files, and terminate the instance using the AWS SDK.
B. Create an AWS Systems Manager document with a script to copy log files to Amazon S3. Create an Auto Scaling lifecycle hook and an Amazon EventBridge (Amazon CloudWatch Events) rule to detect lifecycle events from the Auto Scaling group.
Invoke an AWS Lambda function on the autoscaling:EC2_INSTANCE_TERMINATING transition to call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send CONTINUE to the Auto Scaling group to terminate the instance.
C. Change the log delivery rate to every 5 minutes. Create a script to copy log files to Amazon S3, and add the script to EC2 instance user data Create an Amazon EventBridge (Amazon CloudWatch Events) rule to detect EC2 instance termination.
Invoke an AWS Lambda function from the EventBridge (CloudWatch Events) rule that uses the AWS CLI to run the userdata script to copy the log files and terminate the instance.
D. Create an AWS Systems Manager document with a script to copy log files to Amazon S3. Create an Auto Scaling lifecycle hook that publishes a message to an Amazon Simple Notification Service (Amazon SNS) topic. From the SNS notification call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send ABANDON to the Auto Scaling group to terminate the instance.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/autoscaling/ec2/userguide/configuring-lifecycle-hook-notifications.html

QUESTION 41
Which status represents a failure state in AWS CloudFormation?

A. ROLLBACK_IN_PROGRESS
B. DELETE_IN_PROGRESS
C. UPDATE_COMPLETE_CLEANUP_IN_PROGRESS
D. REVIEW_IN_PROGRESS

Correct Answer: A
Section: (none)

Explanation

Explanation/Reference:
Explanation:
ROLLBACK_IN_PROGRESS means an ongoing removal of one or more stacks after a failed stack creation or after an explicitly canceled stack creation.
DELETE_IN_PROGRESS means an ongoing removal of one or more stacks.
REVIEW_IN_PROGRESS means an ongoing creation of one or more stacks with an expected StackId but without any templates or resources.
UPDATE_COMPLETE_CLEANUP_IN_PROGRESS means an ongoing removal of old resources for one or more stacks after a successful stack update.
Reference: http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-describing-stacks.html

QUESTION 42
A Solutions Architect wants to make sure that only AWS users or roles with suitable permissions can access a new Amazon API Gateway endpoint. The Solutions Architect wants an end-to-end view of each request to analyze the latency of the request and create service maps.
How can the Solutions Architect design the API Gateway access control and perform request inspections?

A. For the API Gateway method, set the authorization to AWS_IAM. Then, give the IAM user or role execute-api:Invoke permission on the REST API resource. Enable the API caller to sign requests with AWS Signature when accessing the endpoint. Use AWS X-Ray to trace and analyze user requests to API Gateway.
B. For the API Gateway resource, set CORS to enabled and only return the company's domain in Access-Control-Allow- Origin headers. Then, give the IAM user or role execute-api:Invoke permission on the REST API resource. Use Amazon CloudWatch to trace and analyze user requests to API Gateway.
C. Create an AWS Lambda function as the custom authorizer, ask the API client to pass the key and secret when making the call, and then use Lambda to validate the key/secret pair against the IAM system. Use AWS X-Ray to trace and analyze user requests to API Gateway.
D. Create a client certificate for API Gateway. Distribute the certificate to the AWS users and roles that need to access the endpoint. Enable the API caller to pass the client certificate when accessing the endpoint. Use Amazon CloudWatch to trace and analyze user requests to API Gateway.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://docs.aws.amazon.com/apigateway/latest/developerguide/http-api-cors.html

QUESTION 43
An organization is planning to host an application on the AWS VPC. The organization wants dedicated instances. However, an AWS consultant advised the organization not to use dedicated instances with VPC as the design has a few limitations.
Which of the below mentioned statements is not a limitation of dedicated instances with VPC?

A. All instances launched with this VPC will always be dedicated instances and the user cannot use a default tenancy model for them.
B. It does not support the AWS RDS with a dedicated tenancy VPC.

C. The user cannot use Reserved Instances with a dedicated tenancy model.
D. The EBS volume will not be on the same tenant hardware as the EC2 instance though the user has configured dedicated tenancy.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Dedicated instances are Amazon EC2 instances that run in a Virtual Private Cloud (VPC) on hardware that is dedicated to a single customer. The client's dedicated instances are physically isolated at the host hardware level from instances that are not dedicated instances as well as from instances that belong to other AWS accounts. All instances launched with the dedicated tenancy model of VPC will always be dedicated instances. Dedicated tenancy has a limitation that it may not support a few services, such as RDS. Even the EBS will not be on dedicated hardware. However, the user can save some cost as well as reserve some capacity by using a Reserved Instance model with dedicated tenancy.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/dedicated-instance.html

QUESTION 44
A company that is developing a mobile game is making game assets available in two AWS Regions. Game assets are served from a set of Amazon EC2 instances behind an Application Load Balancer (ALB) in each Region. The company requires game assets to be fetched from the closest Region. If game assets become unavailable in the closest Region, they should be fetched from the other Region.
What should a solutions architect do to meet these requirements?

A. Create an Amazon CloudFront distribution. Create an origin group with one origin for each ALB. Set one of the origins as primary.
B. Create an Amazon Route 53 health check for each ALCreate a Route 53 failover routing record pointing to the two ALBs. Set the Evaluate Target Health value to Yes.
C. Create two Amazon CloudFront distributions, each with one ALB as the origin. Create an Amazon Route 53 failover routing record pointing to the two CloudFront distributions. Set the Evaluate Target Health value to Yes.
D. Create an Amazon Route 53 health check for each ALB. Create a Route 53 latency alias record pointing to the two ALBs. Set the Evaluate Target Health value to Yes.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 45

Which of the following statements is NOT correct when working with your AWS Direct Connect connection after it is set up completely?

A. You can manage your AWS Direct Connect connections and view the connection details.
B. You can delete a connection as long as there are no virtual interfaces attached to it.
C. You cannot view the current connection ID and verify if it matches the connection ID on the Letter of Authorization (LOA).
D. You can accept a host connection by purchasing a hosted connection from the partner (APN).

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
You can manage your AWS Direct Connect connections and view connection details, accept hosted connections, and delete connections. You can view the current status of your connection. You can also view your connection ID, which looks similar to this example dxcon-xxxx, and verify that it matches the connection ID on the Letter of Authorization (LOA) that you received from Amazon.
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/viewdetails.html

QUESTION 46
A company is migrating an application to the AWS Cloud. The application runs in an on-premises data center and writes thousands of images into a mounted NFS file system each night. After the company migrates the application, the company will host the application on an Amazon EC2 instance with a mounted Amazon Elastic File System (Amazon EFS) file system.
The company has established an AWS Direct Connect connection to AWS. Before the migration cutover, a solutions architect must build a process that will replicate the newly created on-premises images to the EFS file system.
What is the MOST operationally efficient way to replicate the images?

A. Configure a periodic process to run the aws s3 sync command from the on-premises file system to Amazon S3. Configure an AWS Lambda function to process event notifications from Amazon S3 and copy the images from Amazon S3 to the EFS file system.
B. Deploy an AWS Storage Gateway file gateway with an NFS mount point. Mount the file gateway file system on the onpremises server. Configure a process to periodically copy the images to the mount point.
C. Deploy an AWS DataSync agent to an on-premises server that has access to the NFS file system. Send data over the Direct Connect connection to an S3 bucket by using public VIF. Configure an AWS Lambda function to process event notifications from Amazon S3 and copy the images from Amazon S3 to the EFS file system.
D. Deploy an AWS DataSync agent to an on-premises server that has access to the NFS file system. Send data over the Direct Connect connection to an AWS PrivateLink interface VPC endpoint for Amazon EFS by using a private VIF. Configure a DataSync scheduled task to send the images to the EFS file system every 24 hours.

Correct Answer: C
Section: (none)

Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/premiumsupport/knowledge-center/datasync-transfer-efs-cross-region/

QUESTION 47
A large company in Europe plans to migrate its applications to the AWS Cloud. The company uses multiple AWS accounts for various business groups. A data privacy law requires the company to restrict developers' access to AWS European Regions only.
What should the solutions architect do to meet this requirement with the LEAST amount of management overhead?

A. Create IAM users and IAM groups in each account. Create IAM policies to limit access to non-European Regions. Attach the IAM policies to the IAM groups.
B. Enable AWS Organizations, attach the AWS accounts, and create OUs for European Regions and non-European Regions. Create SCPs to limit access to non-European Regions and attach the policies to the OUs.
C. Set up AWS Single Sign-On and attach AWS accounts. Create permission sets with policies to restrict access to non- European Regions. Create IAM users and IAM groups in each account.
D. Enable AWS Organizations, attach the AWS accounts, and create OUs for European Regions and non-European Regions. Create permission sets with policies to restrict access to non-European Regions. Create IAM users and IAM groups in the primary account.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 48
A Solutions Architect must build a highly available infrastructure for a popular global video game that runs on a mobile phone platform. The application runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The database tier is an Amazon RDS MySQL Multi-AZ instance. The entire application stack is deployed in both us-east-1 and eu-central-1. Amazon Route 53 is used to route traffic to the two installations using a latency-based routing policy. A weighted routing policy is configured in Route 53 as a fail over to another region in case the installation in a region becomes unresponsive.
During the testing of disaster recovery scenarios, after blocking access to the Amazon RDS MySQL instance in eu-central-1 from all the application instances running in that region. Route 53 does not automatically failover all traffic to us-east-1.
Based on this situation, which changes would allow the infrastructure to failover to us-east-1? (Choose two.)

A. Specify a weight of 100 for the record pointing to the primary Application Load Balancer in us-east-1 and a weight of 60 for the pointing to the primary Application Load Balancer in eu-central-1.
B. Specify a weight of 100 for the record pointing to the primary Application Load Balancer in us-east-1 and a weight of 0 for the record pointing to the primary Application Load Balancer in eu-central-1.

C. Set the value of Evaluate Target Health to Yes on the latency alias resources for both eu-central-1 and us-east-1.
D. Write a URL in the application that performs a health check on the database layer. Add it as a health check within the weighted routing policy in both regions.
E. Disable any existing health checks for the resources in the policies and set a weight of 0 for the records pointing to primary in both eu-central-1 and us-east-1, and set a weight of 100 for the primary Application Load Balancer only in the region that has healthy resources.

Correct Answer: BC
Section: (none)
Explanation

Explanation/Reference:

QUESTION 49
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 facing 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: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 50
A company has application services that have been containerized and deployed on multiple Amazon EC2 instances with public IPs. An Apache Kafka cluster has been deployed to the EC2 instances. A PostgreSQL database has been migrated to Amazon RDS for PostgreSQL. The company expects a significant increase of orders on its platform when a new version of its flagship product is released.
What changes to the current architecture will reduce operational overhead and support the product release?

A. Create an EC2 Auto Scaling group behind an Application Load Balancer. Create additional read replicas for the DB instance. Create Amazon Kinesis data streams and configure the application services to use the data streams. Store and serve static content directly from Amazon S3.

B. Create an EC2 Auto Scaling group behind an Application Load Balancer. Deploy the DB instance in Multi-AZ mode and enable storage auto scaling. Create Amazon Kinesis data streams and configure the application services to use the data streams. Store and serve static content directly from Amazon S3.

C. Deploy the application on a Kubernetes cluster created on the EC2 instances behind an Application Load Balancer.
Deploy the DB instance in Multi-AZ mode and enable storage auto scaling. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluster. Store static content in Amazon S3 behind an Amazon CloudFront distribution.

D. Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate and enable auto scaling behind an Application Load Balancer. Create additional read replicas for the DB instance. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluster. Store static content in Amazon S3 behind an Amazon CloudFront distribution.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 51
Which of the following Amazon RDS storage types is ideal for applications with light or burst I/O requirements?

A. Both magnetic and Provisioned IOPS storage

B. Magnetic storage

C. Provisioned IOPS storage

D. None of these

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Amazon RDS provides three storage types: magnetic, General Purpose (SSD), and Provisioned IOPS (input/output operations per second). Magnetic (Standard) storage is ideal for applications with light or burst I/O requirements.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Storage.html

QUESTION 52
A company is running a high-user-volume media-sharing application on premises. It currently hosts about 400 TB of data with millions of video files. The company is migrating this application to AWS to improve reliability and reduce costs.

The Solutions Architecture team plans to store the videos in an Amazon S3 bucket and use Amazon CloudFront to distribute videos to users. The company needs to migrate this application to AWS within 10 days with the least amount of downtime possible. The company currently has 1 Gbps connectivity to the Internet with 30 percent free capacity.
Which of the following solutions would enable the company to migrate the workload to AWS and meet all of the requirements?

A. Use a multi-part upload in Amazon S3 client to parallel-upload the data to the Amazon S3 bucket over the Internet. Use the throttling feature to ensure that the Amazon S3 client does not use more than 30 percent of available Internet capacity.
B. Request an AWS Snowmobile with 1 PB capacity to be delivered to the data center. Load the data into Snowmobile and send it back to have AWS download that data to the Amazon S3 bucket. Sync the new data that was generated while migration was in flight.
C. Use an Amazon S3 client to transfer data from the data center to the Amazon S3 bucket over the Internet. Use the throttling feature to ensure the Amazon S3 client does not use more than 30 percent of available Internet capacity.
D. Request multiple AWS Snowball devices to be delivered to the data center. Load the data concurrently into these devices and send it back. Have AWS download that data to the Amazon S3 bucket. Sync the new data that was generated while migration was in flight.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://www.edureka.co/blog/aws-snowball-and-snowmobile-tutorial/

QUESTION 53
A company has an existing on-premises three-tier web application. The Linux web servers serve content from a centralized file share on a NAS server because the content is refreshed several times a day from various sources. The existing infrastructure is not optimized and the company would like to move to AWS in order to gain the ability to scale resources up and down in response to load. On-premises and AWS resources are connected using AWS Direct Connect.
How can the company migrate the web infrastructure to AWS without delaying the content refresh process?

A. Create a cluster of web server Amazon EC2 instances behind a Classic Load Balancer on AWS. Share an Amazon EBS volume among all instances for the content. Schedule a periodic synchronization of this volume and the NAS server.
B. Create an on-premises file gateway using AWS Storage Gateway to replace the NAS server and replicate content to AWS. On the AWS side, mount the same Storage Gateway bucket to each web server Amazon EC2 instance to serve the content.
C. Expose an Amazon EFS share to on-premises users to serve as the NAS serve. Mount the same EFS share to the web server Amazon EC2 instances to serve the content.
D. Create web server Amazon EC2 instances on AWS in an Auto Scaling group. Configure a nightly process where the web server instances are updated from the NAS server.

Correct Answer: B

Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://docs.aws.amazon.com/storagegateway/latest/userguide/GettingStartedAccessFileShare.html

QUESTION 54
A company uses a load balancer to distribute traffic to Amazon EC2 instances in a single Availability Zone. The company is concerned about security and wants a solutions architect to re-architect the solution to meet the following requirements:
Inbound requests must be filtered for common vulnerability attacks.
Rejected requests must be sent to a third-party auditing application. All resources should be highly available.
Which solution meets these requirements?

A. Configure a Multi-AZ Auto Scaling group using the application's AMI. Create an Application Load Balancer (ALB) and select the previously created Auto Scaling group as the target. Use Amazon Inspector to monitor traffic to the ALB and EC2 instances. Create a web ACL in WAF. Create an AWS WAF using the web ACL and ALB. Use an AWS Lambda function to frequently push the Amazon Inspector report to the third-party auditing application
B. Configure an Application Load Balancer (ALB) and add the EC2 instances as targets. Create a web ACL in WAF. Create an AWS WAF using the web ACL and ALB name and enable logging with Amazon CloudWatch Logs. Use an AWS Lambda function to frequently push the logs to the third-party auditing application.
C. Configure an Application Load Balancer (ALB) along with a target group adding the EC2 instances as targets. Create an Amazon Kinesis Data Firehose with the destination of the third-party auditing application. Create a web ACL in WAF. Create an AWS WAF using the web ACL and ALB then enable logging by selecting the Kinesis Data Firehose as the destination.
Subscribe to AWS Managed Rules in AWS Marketplace, choosing the WAF as the subscriber.
D. Configure a Multi-AZ Auto Scaling group using the application's AMI Create an Application Load Balancer (ALB) and select the previously created Auto Scaling group as the target. Create an Amazon Kinesis Data Firehose with a destination of the third-party auditing application. Create a web ACL in WAF. Create an AWS WAF using the web ACL and ALB then enable logging by selecting the Kinesis Data Firehose as the destination. Subscribe to AWS Managed Rules in AWS Marketplace, choosing the WAF as the subscriber.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 55
A company has been using a third-party provider for its content delivery network and recently decided to switch to Amazon CloudFront. The development team wants to maximize performance for the global user base. The company uses a content management system (CMS) that serves both static and dynamic content.

The CMS is behind an Application Load Balancer (ALB) which is set as the default origin for the distribution. Static assets are served from an Amazon S3 bucket. The Origin Access Identity (OAI) was created properly and the S3 bucket policy has been updated to allow the GetObject action from the OAI, but static assets are receiving a 404 error.
Which combination of steps should the solutions architect take to fix the error? (Choose two.)

A. Add another origin to the CloudFront distribution for the static assets.
B. Add a path-based rule to the ALB to forward requests for the static assets.
C. Add an RTMP distribution to allow caching of both static and dynamic content.
D. Add a behavior to the CloudFront distribution for the path pattern and the origin of the static assets.
E. Add a host header condition to the ALB listener and forward the header from CloudFront to add traffic to the allow list.

Correct Answer: AB
Section: (none)
Explanation

Explanation/Reference:

QUESTION 56
A company wants to launch an online shopping website in multiple countries and must ensure that customers are protected against potential “man-in-the-middle” attacks.
Which architecture will provide the MOST secure site access?

A. Use Amazon Route 53 for domain registration and DNS services. Enable DNSSEC for all Route 53 requests. Use AWS Certificate Manager (ACM) to register TLS/SSL certificates for the shopping website, and use Application Load Balancers configured with those TLS/SSL certificates for the site. Use the Server Name Identification extension in all client requests to the site.
B. Register 2048-bit encryption keys from a third-party certificate service. Use a third-party DNS provider that uses the customer managed keys for DNSSec. Upload the keys to ACM, and use ACM to automatically deploy the certificates for secure web services to an EC2 front-end web server fleet by using NGINX. Use the Server Name Identification extension in all client requests to the site.
C. Use Route 53 for domain registration. Register 2048-bit encryption keys from a third-party certificate service. Use a thirdparty DNS service that supports DNSSEC for DNS requests that use the customer managed keys. Import the customer managed keys to ACM to deploy the certificates to Classic Load Balancers configured with those TLS/SSL certificates for the site. Use the Server Name Identification extension in all clients requests to the site.
D. Use Route 53 for domain registration, and host the company DNS root servers on Amazon EC2 instances running Bind.
Enable DNSSEC for DNS requests. Use ACM to register TLS/SSL certificates for the shopping website, and use Application Load Balancers configured with those TLS/SSL certificates for the site. Use the Server Name Identification extension in all client requests to the site.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 57
A large company is migrating its entire IT portfolio to AWS. Each business unit in the company has a standalone AWS account that supports both development and test environments. New accounts to support production workloads will be needed soon.
The Finance department requires a centralized method for payment but must maintain visibility into each group's spending to allocate costs.
The Security team requires a centralized mechanism to control IAM usage in all the company's accounts.
What combination of the following options meet the company's needs with the LEAST effort? (Choose two.)

A. Use a collection of parameterized AWS CloudFormation templates defining common IAM permissions that are launched into each account. Require all new and existing accounts to launch the appropriate stacks to enforce the least privilege model.
B. Use AWS Organizations to create a new organization from a chosen payer account and define an organizational unit hierarchy. Invite the existing accounts to join the organization and create new accounts using Organizations.
C. Require each business unit to use its own AWS accounts. Tag each AWS account appropriately and enable Cost Explorer to administer chargebacks.
D. Enable all features of AWS Organizations and establish appropriate service control policies that filter IAM permissions for sub-accounts.
E. Consolidate all of the company's AWS accounts into a single AWS account. Use tags for billing purposes and IAM's Access Advisor feature to enforce the least privilege model.

Correct Answer: BD
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/ce-what-is.html

QUESTION 58
A company is running a legacy application on Amazon EC2 instances in multiple Availability Zones behind a software load balancer that runs on an active/ standby set of EC2 instances. For disaster recovery, the company has created a warm standby version of the application environment that is deployed in another AWS Region. The domain for the application uses a hosted zone from Amazon Route 53.
The company needs the application to use static IP addresses, even in the case of a failover event to the secondary Region.
The company also requires the client's source IP address to be available for auditing purposes.
Which solution meets these requirements with the LEAST amount of operational overhead?

A. Replace the software load balancer with an AWS Application Load Balancer. Create an AWS Global Accelerator accelerator. Add an endpoint group for each Region. Configure Route 53 health checks. Add an alias record that points to the accelerator.
B. Replace the software load balancer with an AWS Network Load Balancer. Create an AWS Global Accelerator accelerator.

Add an endpoint group for each Region. Configure Route 53 health checks. Add a CNAME record that points to the DNS name of the accelerator.

C. Replace the software load balancer with an AWS Application Load Balancer. Use AWS Global Accelerator to create two separate accelerators. Add an endpoint group for each Region. Configure Route 53 health checks. Add a record set that is configured for active-passive DNS failover. Point the record set to the DNS names of the two accelerators.

D. Replace the software load balancer with an AWS Network Load Balancer. Use AWS Global Accelerator to create two separate accelerators. Add an endpoint group for each Region. Configure Route 53 health checks. Add a record set that is configured for weighted round-robin DNS failover. Point the record set to the DNS names of the two accelerators.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 59
A media company is hosting a high-traffic news website on AWS. The website's front end is based solely on HTML and JavaScript. The company loads all dynamic content by using dynamic asynchronous JavaScript requests to a dedicated backend infrastructure.
The front end runs on four Amazon EC2 instances as web servers. The dynamic backend runs in containers on an Amazon Elastic Container Service (Amazon ECS) cluster that uses an Auto Scaling group of EC2 instances. The ECS tasks are behind an Application Load Balancer (ALB).
Which solutions should a solutions architect recommend to optimize costs? (Choose two.)

A. Migrate the front end of the website to an Amazon S3 bucket Deploy an Amazon CloudFront distribution. Set the S3 bucket as the distribution's origin.

B. Deploy an Amazon CloudFront distribution. Configure the distribution to use the ALB endpoint as the origin.

C. Migrate the front-end services to the ECS cluster. Increase the minimum number of nodes in the Auto Scaling group.

D. Turn on Auto Scaling for the front-end EC2 instances. Configure a new listener rule on the ALB to serve the front end.

E. Migrate the backend of the website to an Amazon S3 bucket. Deploy an Amazon CloudFront distribution. Set the S3 bucket as the distribution's origin.

Correct Answer: BD
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/blogs/security/how-to-enhance-amazon-cloudfront-origin-security-with-aws-waf-andaws-secrets-manager/ https://digitalcloud.training/certification-training/aws-solutions-architect-associate/compute/elasticload-balancing/

QUESTION 60
After setting an AWS Direct Connect, which of the following cannot be done with an AWS Direct Connect Virtual Interface?

A. You can exchange traffic between the two ports in the same region connecting to different Virtual Private Gateways (VGWs) if you have more than one virtual interface.
B. You can change the region of your virtual interface.
C. You can delete a virtual interface; if its connection has no other virtual interfaces, you can delete the connection.
D. You can create a hosted virtual interface.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
You must create a virtual interface to begin using your AWS Direct Connect connection. You can create a public virtual interface to connect to public resources or a private virtual interface to connect to your VPC. Also, it is possible to configure multiple virtual interfaces on a single AWS Direct Connect connection, and you'll need one private virtual interface for each VPC to connect to. Each virtual interface needs a VLAN ID, interface IP address, ASN, and BGP key. To use your AWS Direct Connect connection with another AWS account, you can create a hosted virtual interface for that account. These hosted virtual interfaces work the same as standard virtual interfaces and can connect to public resources or a VPC.
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/WorkingWithVirtualInterfaces.html

QUESTION 61
A medical company is running an application in the AWS Cloud. The application simulates the effect of medical drugs in development.
The application consists of two parts: configuration and simulation. The configuration part runs in AWS Fargate containers in an Amazon Elastic Container Service (Amazon ECS) cluster. The simulation part runs on large, compute optimized Amazon EC2 instances. Simulations can restart if they are interrupted.
The configuration part runs 24 hours a day with a steady load. The simulation part runs only for a few hours each night with a variable load. The company stores simulation results in Amazon S3, and researchers use the results for 30 days. The company must store simulations for 10 years and must be able to retrieve the simulations within 5 hours.
Which solution meets these requirements MOST cost-effectively?

A. Purchase an EC2 Instance Savings Plan to cover the usage for the configuration part. Run the simulation part by using EC2 Spot Instances. Create an S3 Lifecycle policy to transition objects that are older than 30 days to S3 Intelligent-Tiering.
B. Purchase an EC2 Instance Savings Plan to cover the usage for the configuration part and the simulation part. Create an S3 Lifecycle policy to transition objects that are older than 30 days to S3 Glacier.
C. Purchase Compute Savings Plans to cover the usage for the configuration part. Run the simulation part by using EC2 Spot Instances. Create an S3 Lifecycle policy to transition objects that are older than 30 days to S3 Glacier.
D. Purchase Compute Savings Plans to cover the usage for the configuration part. Purchase EC2 Reserved Instances for the simulation part. Create an S3 Lifecycle policy to transition objects that are older than 30 days to S3 Glacier Deep Archive.

Correct Answer: D

Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/savingsplans/faq/

QUESTION 62
A company is currently running a production workload on AWS that is very I/O intensive. Its workload consists of a single tier with 10 c4.8xlarge instances, each with 2 TB gp2 volumes. The number of processing jobs has recently increased, and latency has increased as well. The team realizes that they are constrained on the IOPS. For the application to perform efficiently, they need to increase the IOPS by 3,000 for each of the instances.
Which of the following designs will meet the performance goal MOST cost effectively?

A. Change the type of Amazon EBS volume from gp2 to io1 and set provisioned IOPS to 9,000.
B. Increase the size of the gp2 volumes in each instance to 3 TB.
C. Create a new Amazon EFS file system and move all the data to this new file system. Mount this file system to all 10 instances.
D. Create a new Amazon S3 bucket and move all the data to this new bucket. Allow each instance to access this S3 bucket and use it for storage.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Storage.html

QUESTION 63
A company that tracks medical devices in hospitals wants to migrate its existing storage solution to the AWS Cloud. The company equips all of its devices with sensors that collect location and usage information. This sensor data is sent in unpredictable patterns with large spikes. The data is stored in a MySQL database running on premises at each hospital. The company wants the cloud storage solution to scale with usage.
The company's analytics team uses the sensor data to calculate usage by device type and hospital. The team needs to keep analysis tools running locally while fetching data from the cloud. The team also needs to use existing Java application and SQL queries with as few changes as possible.
How should a solutions architect meet these requirements while ensuring the sensor data is secure?

A. Store the data in an Amazon Aurora Serverless database. Serve the data through a Network Load Balancer (NLB). Authenticate users using the NLB with credentials stored in AWS Secrets Manager.
B. Store the data in an Amazon S3 bucket. Serve the data through Amazon QuickSight using an IAM user authorized with AWS Identity and Access Management (IAM) with the S3 bucket as the data source.
C. Store the data in an Amazon Aurora Serverless database. Serve the data through the Aurora Data API using an IAM user authorized with AWS Identity and Access Management (IAM) and the AWS Secrets Manager ARN.

D. Store the data in an Amazon S3 bucket. Serve the data through Amazon Athena using AWS PrivateLink to secure the data in transit.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 64
A company has an application. Once a month, the application creates a compressed file that contains every object within an Amazon S3 bucket. The total size of the objects before compression is 1 TB.
The application runs by using a scheduled cron job on an Amazon EC2 instance that has a 5 TB Amazon Elastic Block Store (Amazon EBS) volume attached. The application downloads all the files from the source S3 bucket to the EBS volume, compresses the file, and uploads the file to a target S3 bucket. Every invocation of the application takes 2 hours from start to finish.
Which combination of actions should a solutions architect take to OPTIMIZE costs for this application? (Choose two.)

A. Migrate the application to run an AWS Lambda function. Use Amazon EventBridge (Amazon CloudWatch Events) to schedule the Lambda function to run once each month.
B. Configure the application to download the source files by using streams. Direct the streams into a compression library. Direct the output of the compression library into a target object in Amazon S3.
C. Configure the application to download the source files from Amazon S3 and save the files to local storage. Compress the files and upload them to Amazon S3.
D. Configure the application to run as a container in AWS Fargate. Use Amazon EventBridge (Amazon CloudWatch Events) to schedule the task to run once each month.
E. Provision an Amazon Elastic File System (Amazon EFS) file system. Attach the file system to the AWS Lambda function.

Correct Answer: AB
Section: (none)
Explanation

Explanation/Reference:

QUESTION 65
A company has an internal AWS Elastic Beanstalk worker environment inside a VPC that must access an external payment gateway API available on an HTTPS endpoint on the public internet. Because of security policies, the payment gateway's Application team can grant access to only one public IP address.
Which architecture will set up an Elastic Beanstalk environment to access the company's application without making multiple changes on the company's end?

A. Configure the Elastic Beanstalk application to place Amazon EC2 instances in a private subnet with an outbound route to a NAT gateway in a public subnet.

Associate an Elastic IP address to the NAT gateway that can be whitelisted on the payment gateway application side.

B. Configure the Elastic Beanstalk application to place Amazon EC2 instances in a public subnet with an internet gateway. Associate an Elastic IP address to the internet gateway that can be whitelisted on the payment gateway application side.

C. Configure the Elastic Beanstalk application to place Amazon EC2 instances in a private subnet. Set an HTTPS_PROXYapplication parameter to send outbound HTTPS connections to an EC2 proxy server deployed in a public subnet. Associatean Elastic IP address to the EC2 proxy host that can be whitelisted on the payment gateway application side.

D. Configure the Elastic Beanstalk application to place Amazon EC2 instances in a public subnet. Set the HTTPS_PROXYand NO_PROXY application parameters to send non-VPC outbound HTTPS connections to an EC2 proxy server deployedin a public subnet. Associate an Elastic IP address to the EC2 proxy host that can be whitelisted on the payment gatewayapplication side.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/vpc.html

QUESTION 66
A solutions architect needs to implement a client-side encryption mechanism for objects that will be stored in a new Amazon S3 bucket. The solutions architect created a CMK that is stored in AWS Key Management Service (AWS KMS) for this purpose.
The solutions architect created the following IAM policy and attached it to an IAM role:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "DownloadUpload",
            "Action": [
                "s3:GetObject",
                "s3:GetObjectVersion",
                "s3:PutObject",
                "s3:PutObjectAcl"
            ],
            "Effect": "Allow",
            "Resource": "arn:aws:s3:::BucketName/*"
        },
        {
            "Sid": "KMSAccess",
            "Action": [
                "kms:Decrypt",
                "kms:Encrypt"
            ],
            "Effect": "Allow",
            "Resource": "arn:aws:kms:Region:Account:key/Key ID"
        }
    ]
}
```

During tests, the solutions architect was able to successfully get existing test objects in the S3 bucket. However, attempts to upload a new object resulted in an error message. The error message stated that the action was forbidden.
Which action must the solutions architect add to the IAM policy to meet all the requirements?

A. kms:GenerateDataKey
B. kms:GetKeyPolicy
C. kms:GetPublicKey
D. kms:Sign

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 67
An organization is setting up their website on AWS. The organization is working on various security measures to be performed on the AWS EC2 instances. Which of the below mentioned security mechanisms will not help the organization to avoid future data leaks and identify security weaknesses?

A. Run penetration testing on AWS with prior approval from Amazon.
B. Perform SQL injection for application testing.
C. Perform a Code Check for any memory leaks.
D. Perform a hardening test on the AWS instance.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
AWS security follows the shared security model where the user is as much responsible as Amazon. Since Amazon is a public cloud it is bound to be targeted by hackers. If an organization is planning to host their application on AWS EC2, they should perform the below mentioned security checks as a measure to find any security weakness/data leaks:
Perform penetration testing as performed by attackers to find any vulnerability. The organization must take an approval from AWS before performing penetration testing Perform hardening testing to find if there are any unnecessary ports open Perform SQL injection to find any DB security issues The code memory checks are generally useful when the organization wants to improve the application performance.
Reference: http://aws.amazon.com/security/penetration-testing/

QUESTION 68
When does an AWS Data Pipeline terminate the AWS Data Pipeline-managed compute resources?

A. AWS Data Pipeline terminates AWS Data Pipeline-managed compute resources every 2 hours.
B. When the final activity that uses the resources is running
C. AWS Data Pipeline terminates AWS Data Pipeline-managed compute resources every 12 hours.
D. When the final activity that uses the resources has completed successfully or failed

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Compute resources will be provisioned by AWS Data Pipeline when the first activity for a scheduled time that uses those resources is ready to run, and those

instances will be terminated when the final activity that uses the resources has completed successfully or failed.
Reference:
https://aws.amazon.com/datapipeline/faqs/

QUESTION 69
You have just added a new instance to your Auto Scaling group, which receives ELB health checks. An ELB heath check says the new instance's state is out of Service.
What does Auto Scaling do in this particular scenario?

A. It replaces the instance with a healthy one
B. It stops the instance
C. It marks an instance as unhealthy
D. It terminates the instance

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
If you have attached a load balancer to your Auto Scaling group, you can have Auto Scaling include the results of Elastic Load Balancing health checks when it determines the health status of an instance. After you add ELB health checks, Auto Scaling will mark an instance as unhealthy if Elastic Load Balancing reports the instance state as Out of Service. Frequently, an Auto Scaling instance that has just come into service needs to warm up before it can pass the Auto Scaling health check.
Auto Scaling waits until the health check grace period ends before checking the health status of the instance. While the EC2 status checks and ELB health checks can complete before the health check grace period expires, Auto Scaling does not act on them until the health check grace period expires. To provide ample warm-up time for your instances, ensure that the health check grace period covers the expected startup time for your application.
Reference: http://docs.aws.amazon.com/autoscaling/latest/userguide/healthcheck.html

QUESTION 70
A company has an application that uses Amazon EC2 instances in an Auto Scaling group. The Quality Assurance (QA) department needs to launch a large number of short-lived environments to test the application. The application environments are currently launched by the Manager of the department using an AWS CloudFormation template. To launch the stack, the Manager uses a role with permission to use CloudFormation, EC2, and Auto Scaling APIs. The Manager wants to allow testers to launch their own environments, but does not want to grant broad permissions to each user.
Which set up would achieve these goals?

A. Upload the AWS CloudFormation template to Amazon S3. Give users in the QA department permission to assume the Manager's role and add a policy that restricts the permissions to the template and the resources it creates. Train users to launch the template from the CloudFormation console.
B. Create an AWS Service Catalog product from the environment template. Add a launch constraint to the product with the existing role. Give users in the QA department permission to use AWS Service Catalog APIs only. Train users to launch the templates from the AWS Service Catalog console.

C. Upload the AWS CloudFormation template to Amazon S3. Give users in the QA department permission to use CloudFormation and S3 APIs, with conditions that restrict the permission to the template and the resources it creates. Train users to launch the template from the CloudFormation console.
D. Create an AWS Elastic Beanstalk application from the environment template. Give users in the QA department permission to use Elastic Beanstalk permissions only. Train users to launch Elastic Beanstalk environment with the Elastic Beanstalk CLI, passing the existing role to the environment as a service role.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://aws.amazon.com/ru/blogs/mt/how-to-launch-secure-and-governed-aws-resources-with-aws-cloudformation-and-awsservice-catalog/

QUESTION 71
A company is using AWS to run an internet-facing production application written in Node.js. The Development team is responsible for pushing new versions of their software directly to production. The application software is updated multiple times a day. The team needs guidance from a Solutions Architect to help them deploy the software to the production fleet quickly and with the least amount of disruption to the service.
Which option meets these requirements?

A. Prepackage the software into an AMI and then use Auto Scaling to deploy the production fleet. For software changes, update the AMI and allow Auto Scaling to automatically push the new AMI to production.
B. Use AWS CodeDeploy to push the prepackaged AMI to production. For software changes, reconfigure CodeDeploy with new AMI identification to push the new AMI to the production fleet.
C. Use AWS Elastic Beanstalk to host the production application. For software changes, upload the new application version to Elastic Beanstalk to push this to the production fleet using a blue/green deployment method.
D. Deploy the base AMI through Auto Scaling and bootstrap the software using user data. For software changes, SSH to each of the instances and replace the software with the new version.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 72
A company runs a dynamic mission-critical web application that has an SLA of 99.99%. Global application users access the application 24/7. The application is currently hosted on premises and routinely fails to meet its SLA, especially when millions of users access the application concurrently. Remote users complain of

latency.
How should this application be redesigned to be scalable and allow for automatic failover at the lowest cost?

A. Use Amazon Route 53 failover routing with geolocation-based routing. Host the website on automatically scaled Amazon EC2 instances behind an Application Load Balancer with an additional Application Load Balancer and EC2 instances for the application layer in each region. Use a Multi-AZ deployment with MySQL as the data layer.
B. Use Amazon Route 53 round robin routing to distribute the load evenly to several regions with health checks. Host the website on automatically scaled Amazon ECS with AWS Fargate technology containers behind a Network Load Balancer, with an additional Network Load Balancer and Fargate containers for the application layer in each region. Use Amazon Aurora replicas for the data layer.
C. Use Amazon Route 53 latency-based routing to route to the nearest region with health checks. Host the website in Amazon S3 in each region and use Amazon API Gateway with AWS Lambda for the application layer. Use Amazon DynamoDB global tables as the data layer with Amazon DynamoDB Accelerator (DAX) for caching.
D. Use Amazon Route 53 geolocation-based routing. Host the website on automatically scaled AWS Fargate containers behind a Network Load Balancer with an additional Network Load Balancer and Fargate containers for the application layer in each region. Use Amazon Aurora Multi-Master for Aurora MySQL as the data layer.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://aws.amazon.com/getting-started/hands-on/build-serverless-web-app-lambda-apigateway-s3-dynamodbcognito/module-3/

QUESTION 73
An organization is making software for the CIA in USA. CIA agreed to host the application on AWS but in a secure environment. The organization is thinking of hosting the application on the AWS GovCloud region. Which of the below mentioned difference is not correct when the organization is hosting on the AWS GovCloud in comparison with the AWS standard region?

A. The billing for the AWS GovCLoud will be in a different account than the Standard AWS account.
B. GovCloud region authentication is isolated from Amazon.com.
C. Physical and logical administrative access only to U.S. persons.
D. It is physically isolated and has logical network isolation from all the other regions.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
AWS GovCloud (US) is an isolated AWS region designed to allow U.S. government agencies and customers to move sensitive workloads into the cloud by addressing their specific regulatory and compliance requirements. The AWS GovCloud (US) Region adheres to the U.S. International Traffic in Arms Regulations (ITAR) requirements. It has added advantages, such as:
Restricting physical and logical administrative access to U.S. persons only There will be a separate AWS GovCloud (US) credentials, such as access key and secret access key than the standard AWS account The user signs in with the IAM user name and password The AWS GovCloud (US) Region authentication is completely isolated from Amazon.com If the organization is planning to host on EC2 in AWS GovCloud then it will be billed to standard AWS account of organization since AWS GovCloud billing is linked with the standard AWS account and is not be billed separately.
Reference: http://docs.aws.amazon.com/govcloud-us/latest/UserGuide/whatis.html

QUESTION 74
A company is storing data on premises on a Windows file server. The company produces 5 GB of new data daily.
The company migrated part of its Windows-based workload to AWS and needs the data to be available on a file system in the cloud. The company already has established an AWS Direct Connect connection between the on-premises network and AWS. Which data migration strategy should the company use?

A. Use the file gateway option in AWS Storage Gateway to replace the existing Windows file server, and point the existing file share to the new file gateway
B. Use AWS DataSync to schedule a daily task to replicate data between the on-premises Windows file server and Amazon FSx
C. Use AWS Data Pipeline to schedule a daily task to replicate data between the on-premises Windows file server and Amazon Elastic File System (Amazon EFS)
D. Use AWS DataSync to schedule a daily task to replicate data between the on-premises Windows file server and Amazon Elastic File System (Amazon EFS)

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 75
A company standardized its method of deploying applications to AWS using AWS CodePipeline and AWS Cloud Formation.
The applications are in TypeScript and Python. The company has recently acquired another business that deploys applications to AWS using Python scripts.
Developers from the newly acquired company are hesitant to move their applications under Cloud Formation because it would require that they learn a new domain-specific language and eliminate their access to language features, such as looping.
How can the acquired applications quickly be brought up to deployment standards while addressing the developers' concerns?

A. Create Cloud Formation templates and re-use parts of the Python scripts as Instance user data. Use the AWS Cloud Development Kit (AWS CDK) to deploy the application using these templates. Incorporate the AWS CDK into CodePipeline and deploy the application to AWS using these templates.
B. Use a third-party resource provisioning engine inside AWS CodeBuild to standardize the deployment processes of the existing and acquired company. Orchestrate the CodeBuild job using CodePipeline.

C. Standardize on AWS OpsWorks. Integrate OpsWorks with CodePipeline. Have the developers create Chef recipes to deploy their applications on AWS.
D. Define the AWS resources using TypeScript or Python. Use the AWS Cloud Development Kit (AWS CDK) to create CloudFormation templates from the developers' code, and use the AWS CDK to create CloudFormation stacks. Incorporate the AWS CDK as a CodeBuild job in CodePipeline.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 76
A bucket owner has allowed another account's IAM users to upload or access objects in his bucket. The IAM user of Account A is trying to access an object created by the IAM user of account
B. What will happen in this scenario?

A. It is not possible to give permission to multiple IAM users
B. What will happen in this scenario?
AWS S3 will verify proper rights given by the owner of Account A, the bucket owner as well as by the IAM user B to the object
C. The bucket policy may not be created as S3 will give error due to conflict of Access Rights
D. It is not possible that the IAM user of one account accesses objects of the other IAM user

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
If a IAM user is trying to perform some action on an object belonging to another AWS user's bucket, S3 will verify whether the owner of the IAM user has given sufficient permission to him. It also verifies the policy for the bucket as well as the policy defined by the object owner.
Reference: http://docs.aws.amazon.com/AmazonS3/latest/dev/access-control-auth-workflow-object-operation.html

QUESTION 77
A retail company is hosting an ecommerce website on AWS across multiple AWS Regions. The company wants the website to be operational at all times for online purchases. The website stores data in an Amazon RDS for MySQL DB instance.
Which solution will provide the HIGHEST availability for the database?

A. Configure automated backups on Amazon RDS. In the case of disruption, promote an automated backup to be a standalone DB instance. Direct database traffic to the promoted DB instance. Create a replacement read replica that has the promoted DB instance as its source.

B. Configure global tables and read replicas on Amazon RDS. Activate the cross-Region scope. In the case of disruption, use AWS Lambda to copy the read replicas from one Region to another Region.
C. Configure global tables and automated backups on Amazon RDS. In the case of disruption, use AWS Lambda to copy the read replicas from one Region to another Region.
D. Configure read replicas on Amazon RDS. In the case of disruption, promote a cross-Region and read replica to be a standalone DB instance. Direct database traffic to the promoted DB instance. Create a replacement read replica that has the promoted DB instance as its source.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html

QUESTION 78
You have been asked to set up a public website on AWS with the following criteria:
You want the database and the application server running on an Amazon VPC. You want the database to be able to connect to the Internet so that it can be automatically updated to the correct patch level. You do not want to receive any incoming traffic from the Internet to the database.
Which solutions would be the best to satisfy all the above requirements for your planned public website on AWS? (Choose two.)

A. Set up both the public website and the database on a public subnet and block all incoming requests from the Internet with a Network Access Control List (NACL)
B. Set up both the public website and the database on a public subnet, and block all incoming requests from the Internet with a security group which only allows access from the IP of the public website.
C. Set up the public website on a public subnet and set up the database in a private subnet which connects to the Internet via a NAT instance.
D. Set up both the public website and the database on a private subnet and block all incoming requests from the Internet with a Network Access Control List (NACL). Set up a Security group between the public website and the database which only allows access via port 80.

Correct Answer: BC
Section: (none)
Explanation

Explanation/Reference:
Explanation:
You want the database to be able to connect to the Internet you need to either set it up on a public subnet or set it up on a private subnet which connects to the Internet via a NAT instance
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario2.html

QUESTION 79

In regard to DynamoDB, for which one of the following parameters does Amazon not charge you?

A. Storage cost
B. I/O usage within the same Region
C. Cost per provisioned read units
D. Cost per provisioned write units

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In DynamoDB, you will be charged for the storage and the throughput you use rather than for the I/O which has been used.
Reference: http://aws.amazon.com/dynamodb/pricing/

QUESTION 80
A company has a High Performance Computing (HPC) cluster in its on-premises data center, which runs thousands of jobs in parallel for one week every month, processing petabytes of images. The images are stored on a network file server, which is replicated to a disaster recovery site. The on-premises data center has reached capacity and has started to spread the jobs out over the course of the month in order to better utilize the cluster, causing a delay in the job completion. The company has asked its Solutions Architect to design a cost-effective solution on AWS to scale beyond the current capacity of 5,000 cores and 10 petabytes of data. The solution must require the least amount of management overhead and maintain the current level of durability.
Which solution will meet the company's requirements?

A. Create a container in the Amazon Elastic Container Registry with the executable file for the job. Use Amazon ECS with Spot Fleet in Auto Scaling groups. Store the raw data in Amazon EBS SC1 volumes and write the output to Amazon S3.
B. Create an Amazon EMR cluster with a combination of On Demand and Reserved Instance Task Nodes that will use Spark to pull data from Amazon S3. Use Amazon DynamoDB to maintain a list of jobs that need to be processed by the Amazon EMR cluster.
C. Store the raw data in Amazon S3, and use AWS Batch with Managed Compute Environments to create Spot Fleets. Submit jobs to AWS Batch Job Queues to pull down objects from Amazon S3 onto Amazon EBS volumes for temporary storage to be processed, and then write the results back to Amazon S3.
D. Submit the list of jobs to be processed to an Amazon SQS to queue the jobs that need to be processed. Create a diversified cluster of Amazon EC2 worker instances using Spot Fleet that will automatically scale based on the queue depth. Use Amazon EFS to store all the data sharing it across all instances in the cluster.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 81
A solutions architect is implementing federated access to AWS for users of the company's mobile application. Due to regulatory and security requirements, the application must use a custom-built solution for authenticating users and must use IAM roles for authorization.
Which of the following actions would enable authentication and authorization and satisfy the requirements? (Choose two.)

A. Use a custom-built SAML-compatible solution for authentication and AWS SSO for authorization.
B. Create a custom-built LDAP connector using Amazon API Gateway and AWS Lambda for authentication. Store authorization tokens in Amazon DynamoDB, and validate authorization requests using another Lambda function that reads the credentials from DynamoDB.
C. Use a custom-built OpenID Connect-compatible solution with AWS SSO for authentication and authorization.
D. Use a custom-built SAML-compatible solution that uses LDAP for authentication and uses a SAML assertion to perform authorization to the IAM identity provider.
E. Use a custom-built OpenID Connect-compatible solution for authentication and use Amazon Cognito for authorization.

Correct Answer: AC
Section: (none)
Explanation

Explanation/Reference:

QUESTION 82
An organization is hosting a scalable web application using AWS. The organization has configured ELB and Auto Scaling to make the application scalable.
Which of the below mentioned statements is not required to be followed for ELB when the application is planning to host a web application on VPC?

A. The ELB and all the instances should be in the same subnet.
B. Configure the security group rules and network ACLs to allow traffic to be routed between the subnets in the VPC.
C. The internet facing ELB should have a route table associated with the internet gateway.
D. The internet facing ELB should be only in a public subnet.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web

Services (AWS) cloud. The user has complete control over the virtual networking environment.
Within this virtual private cloud, the user can launch AWS resources, such as an ELB, and EC2 instances. There are two ELBs available with VPC: internet facing and internal (private) ELB. For the internet facing ELB it is required that the ELB should be in a public subnet. After the user creates the public subnet, he should ensure to associate the route table of the public subnet with the internet gateway to enable the load balancer in the subnet to connect with the internet. The ELB and instances can be in a separate subnet. However, to allow communication between the instance and the ELB the user must configure the security group rules and network ACLs to allow traffic to be routed between the subnets in his VPC.
Reference: http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/CreateVPCForELB.html

QUESTION 83
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 life cycle 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: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 84
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 1KB 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:
Explanation:
The POC solution is being scaled up by 1000, which means it will require 72TB of Storage to retain 24 months' worth of data.
This rules out RDS as a possible DB solution which leaves you with Redshift. I believe DynamoDB is a more cost effective and scales better for ingest rather than using EC2 in an auto scaling group. Also, this example solution from AWS is somewhat similar for reference.

QUESTION 85
You have subscribed to the AWS Business and Enterprise support plan.
Your business has a backlog of problems, and you need about 20 of your IAM users to open technical support cases.
How many users can open technical support cases under the AWS Business and Enterprise support plan?

A. 5 users
B. 10 users
C. Unlimited
D. 1 user

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In the context of AWS support, the Business and Enterprise support plans allow an unlimited number of users to open technical support cases (supported by AWS Identity and Access Management (IAM)).
Reference: https://aws.amazon.com/premiumsupport/faqs/

QUESTION 86
What RAID method is used on the Cloud Block Storage back-end to implement a very high level of reliability and performance?

A. RAID 1 (Mirror)
B. RAID 5 (Blocks striped, distributed parity)
C. RAID 10 (Blocks mirrored and striped)
D. RAID 2 (Bit level striping)

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Cloud Block Storage back-end storage volumes employs the RAID 10 method to provide a very high level of reliability and performance.
Reference: http://www.rackspace.com/knowledge_center/product-faq/cloud-block-storage

QUESTION 87
A user is planning to launch multiple EC2 instance same as current running instance.
Which of the below mentioned parameters is not copied by Amazon EC2 in the launch wizard when the user has selected the option "Launch more like this"?

A. Termination protection
B. Tenancy setting
C. Storage
D. Shutdown behavior

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The Amazon EC2 console provides a "Launch more like this" wizard option that enables the user to use a current instance as a template for launching other instances. This option automatically populates the Amazon EC2 launch wizard with certain configuration details from the selected instance.
The following configuration details are copied from the selected instance into the launch wizard: AMI ID Instance type Availability Zone, or the VPC and subnet in which the selected instance is located Public IPv4 address. If the selected instance currently has a public IPv4 address, the new instance receives a public IPv4 address - regardless of the selected instance's default public IPv4 address setting.
For more information about public IPv4 addresses, see Public IPv4 Addresses and External DNS Hostnames.
Placement group, if applicable
IAM role associated with the instance, if applicable Shutdown behavior setting (stop or terminate) Termination protection setting (true or false) CloudWatch monitoring (enabled or disabled) Amazon EBS-optimization setting (true or false) Tenancy setting, if launching into a VPC (shared or dedicated) Kernel ID and

RAM disk ID, if applicable User data, if specified Tags associated with the instance, if applicable Security groups associated with the instance The following configuration details are not copied from your selected instance; instead, the wizard applies their default settings or behavior:
(VPC only) Number of network interfaces: The default is one network interface, which is the primary network interface (eth0).
Storage: The default storage configuration is determined by the AMI and the instance type.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/launching-instance.html

QUESTION 88
A gaming company created a game leaderboard by using a Multi-AZ deployment of an Amazon RDS database. The number of users is growing, and the queries to get individual player rankings are getting slower over time. The company expects a surge in users for an upcoming version and wants to optimize the design for scalability and performance.
Which solution will meet these requirements?

A. Migrate the database to Amazon DynamoDB. Store the leaderboard data in two different tables. Use Apache HiveQL JOIN statements to build the leaderboard.
B. Keep the leaderboard data in the RDS DB instance. Provision a Multi-AZ deployment of an Amazon ElastiCache for Redis cluster.
C. Stream the leaderboard data by using Amazon Kinesis Data Firehose with an Amazon S3 bucket as the destination. Query the S3 bucket by using Amazon Athena for the leaderboard.
D. Add a read-only replica to the RDS DB instance. Add an RDS Proxy database proxy.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/rds-proxy.html

QUESTION 89
An IoT company has rolled out a fleet of sensors for monitoring temperatures in remote locations. Each device connects to AWS IoT Core and sends a message 30 seconds, updating an Amazon DynamoDB table. A System Administrator users AWS IoT to verify the devices are still sending messages to AWS IoT Core: the database is not updating.
What should a Solutions Architect check to determine why the database is not being updated?

A. Verify the AWS IoT Device Shadow service is subscribed to the appropriate topic and is executing the AWS Lambda function.
B. Verify that AWS IoT monitoring shows that the appropriate AWS IoT rules are being executed, and that the AWS IoT rules are enabled with the correct rule actions.
C. Check the AWS IoT Fleet indexing service and verify that the thing group has the appropriate IAM role to update DynamoDB.
D. Verify that AWS IoT things are using MQTT instead of MQTT over WebSocket, then check that the provisioning has the appropriate policy attached.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 90
A large company runs workloads in VPCs that are deployed across hundreds of AWS accounts. Each VPC consists of public subnets and private subnets that span across multiple Availability Zones. NAT gateways are deployed in the public subnets and allow outbound connectivity to the internet from the private subnets.
A solutions architect is working on a hub-and-spoke design. All private subnets in the spoke VPCs must route traffic to the internet through an egress VPC. The solutions architect already has deployed a NAT gateway in an egress VPC in a central AWS account.
Which set of additional steps should the solutions architect take to meet these requirements?

A. Create peering connections between the egress VPC and the spoke VPCs. Configure the required routing to allow access to the internet.
B. Create a transit gateway, and share it with the existing AWS accounts. Attach existing VPCs to the transit gateway. Configure the required routing to allow access to the internet.
C. Create a transit gateway in every account. Attach the NAT gateway to the transit gateways. Configure the required routing to allow access to the internet.
D. Create an AWS PrivateLink connection between the egress VPC and the spoke VPCs. Configure the required routing to allow access to the internet.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html

QUESTION 91
A user is trying to create a PIOPS EBS volume with 3 GB size and 90 IOPS. Will AWS create the volume?

A. No, since the PIOPS and EBS size ratio is less than 30
B. Yes, since the ratio between EBS and IOPS is less than 30
C. No, the EBS size is less than 4GB
D. Yes, since PIOPS is higher than 100

Correct Answer: C
Section: (none)

Explanation

Explanation/Reference:
Explanation:
A Provisioned IOPS (SSD) volume can range in size from 4 GiB to 16 TiB and you can provision up to 20,000 IOPS per volume.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html#EBSVolumeTypes_piops

QUESTION 92
A bank is re-architecting its mainframe-based credit card approval processing application to a cloud-native application on the AWS cloud.
The new application will receive up to 1,000 requests per second at peak load. There are multiple steps to each transaction, and each step must receive the result of the previous step. The entire request must return an authorization response within less than 2 seconds with zero data loss. Every request must receive a response. The solution must be Payment Card Industry Data Security Standard (PCI DSS)-compliant.
Which option will meet all of the bank's objectives with the LEAST complexity and LOWEST cost while also meeting compliance requirements?

A. Create an Amazon API Gateway to process inbound requests using a single AWS Lambda task that performs multiple steps and returns a JSON object with the approval status. Open a support case to increase the limit for the number of concurrent Lambdas to allow room for bursts of activity due to the new application.
B. Create an Application Load Balancer with an Amazon ECS cluster on Amazon EC2 Dedicated Instances in a target group to process incoming requests. Use Auto Scaling to scale the cluster out/in based on average CPU utilization. Deploy a web service that processes all of the approval steps and returns a JSON object with the approval status.
C. Deploy the application on Amazon EC2 on Dedicated Instances. Use an Elastic Load Balancer in front of a farm of application servers in an Auto Scaling group to handle incoming requests. Scale out/in based on a custom Amazon CloudWatch metric for the number of inbound requests per second after measuring the capacity of a single instance.
D. Create an Amazon API Gateway to process inbound requests using a series of AWS Lambda processes, each with an Amazon SQS input queue. As each step completes, it writes its result to the next step's queue. The final step returns a JSON object with the approval status. Open a support case to increase the limit for the number of concurrent Lambdas to allow room for bursts of activity due to the new application.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 93
Which of the following components of AWS Data Pipeline polls for tasks and then performs those tasks?

A. Pipeline Definition
B. Task Runner
C. Amazon Elastic MapReduce (EMR)

D. AWS Direct Connect

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Task Runner polls for tasks and then performs those tasks.
Reference: http://docs.aws.amazon.com/datapipeline/latest/DeveloperGuide/what-is-datapipeline.html

QUESTION 94
A solutions architect is designing a network for a new cloud deployment. Each account will need autonomy to modify route tables and make changes. Centralized and controlled egress internet connectivity is also needed. The cloud footprint is expected to grow to thousands of AWS accounts.
Which architecture will meet these requirements?

A. A centralized transit VPC with a VPN connection to a standalone VPC in each account. Outbound internet traffic will be controlled by firewall appliances.
B. A centralized shared VPC with a subnet for each account. Outbound internet traffic will be controlled through a fleet of proxy servers.
C. A shared services VPC to host central assets to include a fleet of firewalls with a route to the internet. Each spoke VPC will peer to the central VPC.
D. A shared transit gateway to which each VPC will be attached. Outbound internet access will route through a fleet of VPNattached firewalls.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 95
A company has an organization that has many AWS accounts in AWS Organizations. A solutions architect must improve how the company manages common security group rules for the AWS accounts in the organization.
The company has a common set of IP CIDR ranges in an allow list in each AWS account to allow access to and from the company's on-premises network. Developers within each account are responsible for adding new IP CIDR ranges to their security groups. The security team has its own AWS account. Currently, the security team notifies the owners of the other AWS accounts when changes are made to the allow list.
The solutions architect must design a solution that distributes the common set of CIDR ranges across all accounts.
Which solution meets these requirements with the LEAST amount of operational overhead?

A. Set up an Amazon Simple Notification Service (Amazon SNS) topic in the security team's AWS account. Deploy an AWS Lambda function in each AWS account. Configure the Lambda function to run every time an SNS topic receives a message.
Configure the Lambda function to take an IP address as input and add it to a list of security groups in the account. Instruct the security team to distribute

changes by publishing messages to its SNS topic.

B. Create new customer-managed prefix lists in each AWS account within the organization. Populate the prefix lists in each account with all internal CIDR ranges. Notify the owner of each AWS account to allow the new customer-managed prefix list IDs in their accounts in their security groups. Instruct the security team to share updates with each AWS account owner.

C. Create a new customer-managed prefix list in the security team's AWS account. Populate the customer-managed prefix list with all internal CIDR ranges. Share the customer-managed prefix list with the organization by using AWS Resource Access Manager. Notify the owner of each AWS account to allow the new customer-managed prefix list ID in their security groups.

D. Create an IAM role in each account in the organization. Grant permissions to update security groups. Deploy an AWS Lambda function in the security team's AWS account. Configure the Lambda function to take a list of internal IP addresses as input, assume a role in each organization account, and add the list of IP addresses to the security groups in each account.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/vpc/latest/userguide/managed-prefix-lists.html

QUESTION 96
A company's site reliability engineer is performing a review of Amazon FSx for Windows File Server deployments within an account that the company acquired. Company policy states that all Amazon FSx file systems must be configured to be highly available across Availability Zones.
During the review, the site reliability engineer discovers that one of the Amazon FSx file systems used a deployment type of Single-AZ 2. A solutions architect needs to minimize downtime while aligning this Amazon FSx file system with company policy.
What should the solutions architect do to meet these requirements?

A. Reconfigure the deployment type to Multi-AZ for this Amazon FSx file system.

B. Create a new Amazon FSx file system with a deployment type of Multi-AZ. Use AWS DataSync to transfer data to the new Amazon FSx file system. Point users to the new location.

C. Create a second Amazon FSx file system with a deployment type of Single-AZ 2. Use AWS DataSync to keep the data in sync. Switch users to the second Amazon FSx file system in the event of failure.

D. Use the AWS Management Console to take a backup of the Amazon FSx file system. Create a new Amazon FSx file system with a deployment type of Multi-AZ. Restore the backup to the new Amazon FSx file system Point users to the new location.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

Explanation:
Reference: https://docs.aws.amazon.com/fsx/latest/WindowsGuide/high-availability-multiAZ.html

QUESTION 97
A company has a legacy monolithic application that is critical to the company's business. The company hosts the application on an Amazon EC2 instance that runs Amazon Linux 2. The company's application team receives a directive from the legal department to back up the data from the instance's encrypted Amazon Elastic Block Store (Amazon EBS) volume to an Amazon S3 bucket. The application team does not have the administrative SSH key pair for the instance. The application must continue to serve the users.
Which solution will meet these requirements?

A. Attach a role to the instance with permission to write to Amazon S3. Use the AWS Systems Manager Session Manager option to gain access to the instance and run commands to copy data into Amazon S3.
B. Create an image of the instance with the reboot option turned on. Launch a new EC2 instance from the image. Attach a role to the new instance with permission to write to Amazon S3. Run a command to copy data into Amazon S3.
C. Take a snapshot of the EBS volume by using Amazon Data Lifecycle Manager (Amazon DLM). Copy the data to Amazon S3.
D. Create an image of the instance. Launch a new EC2 instance from the image. Attach a role to the new instance with permission to write to Amazon S3. Run a command to copy data into Amazon S3.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonS3.html

QUESTION 98
Doug has created a VPC with CIDR 10.201.0.0/16 in his AWS account. In this VPC he has created a public subnet with CIDR block 10.201.31.0/24. While launching a new EC2 from the console, he is not able to assign the private IP address 10.201.31.6 to this instance.
Which is the most likely reason for this issue?

A. Private address IP 10.201.31.6 is currently assigned to another interface
B. Private IP address 10.201.31.6 is reserved by Amazon for IP networking purposes.
C. Private IP address 10.201.31.6 is blocked via ACLs in Amazon infrastructure as a part of platform security.
D. Private IP address 10.201.31.6 is not part of the associated subnet's IP address range.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In Amazon VPC, you can assign any Private IP address to your instance as long as it is: Part of the associated subnet's IP address range Not reserved by Amazon for IP networking purposes Not currently assigned to another interface
Reference: http://aws.amazon.com/vpc/faqs/

QUESTION 99
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:
Explanation:
Here are some of the things that you can build using fine-grained access control:
A mobile app that displays information for nearby airports, based on the user's location. The app can access and display attributes such airline names, arrival times, and flight numbers. However, it cannot access or display pilot names or passenger counts.
A mobile game which stores high scores for all users in a single table. Each user can update their own scores, but has no access to the other ones. Reference: https://aws.amazon.com/blogs/aws/fine-grained-access-control-for-amazon-dynamodb/

QUESTION 100
You are playing around with setting up stacks using JSON templates in CloudFormation to try and understand them a little better. You have set up about 5 or 6 but now start to wonder if you are being charged for these stacks.
What is AWS's billing policy regarding stack resources?

A. You are not charged for the stack resources if they are not taking any traffic.
B. You are charged for the stack resources for the time they were operating (but not if you deleted the stack within 30 minutes)
C. You are charged for the stack resources for the time they were operating (but not if you deleted the stack within 60 minutes)
D. You are charged for the stack resources for the time they were operating (even if you deleted the stack right away)

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A stack is a collection of AWS resources that you can manage as a single unit. In other words, you can create, update, or delete a collection of resources by creating, updating, or deleting stacks. All the resources in a stack are defined by the stack's AWS CloudFormation template. A stack, for instance, can include all the resources required to run a web application, such as a web server, a database, and networking rules. If you no longer require that web application, you can simply delete the stack, and all of its related resources are deleted. You are charged for the stack resources for the time they were operating (even if you deleted the stack right away).
Reference: http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacks.html

QUESTION 101
A company is migrating its marketing website and content management system from an on-premises data center to AWS.
The company wants the AWS application to be deployed in a VPC with Amazon EC2 instances used for the web servers and an Amazon RDS instance for the database.
The company has a runbook document that describes the installation process of the on-premises system. The company would like to base the AWS system on the processes referenced in the runbook document. The runbook document describes the installation and configuration of the operating systems, network settings, the website, and content management system software on the servers. After the migration is complete, the company wants to be able to make changes quickly to take advantage of other AWS features.
How can the application and environment be deployed and automated in AWS, while allowing for future changes?

A. Update the runbook to describe how to create the VPC, the EC2 instances, and the RDS instance for the application by using the AWS Console. Make sure that the rest of the steps in the runbook are updated to reflect any changes that may come from the AWS migration.
B. Write a Python script that uses the AWS API to create the VPC, the EC2 instances, and the RDS instance for the application. Write shell scripts that implement the rest of the steps in the runbook. Have the Python script copy and run the shell scripts on the newly created instances to complete the installation.
C. Write an AWS CloudFormation template that creates the VPC, the EC2 instances, and the RDS instance for the application. Ensure that the rest of the steps in the runbook are updated to reflect any changes that may come from the AWS migration.
D. Write an AWS CloudFormation template that creates the VPC, the EC2 instances, and the RDS instance for the application. Include EC2 user data in the AWS CloudFormation template to install and configure the software.

Correct Answer: D

Section: (none)
Explanation

Explanation/Reference:

QUESTION 102
A company is migrating a subset of its application APIs from Amazon EC2 instances to run on a serverless infrastructure.
The company has set up Amazon API Gateway, AWS Lambda, and Amazon DynamoDB for the new application. The primary responsibility of the Lambda function is to obtain data from a third-party Software as a Service (SaaS) provider. For consistency, the Lambda function is attached to the same virtual private cloud (VPC) as the original EC2 instances.
Test users report an inability to use this newly moved functionality, and the company is receiving 5xx errors from API Gateway. Monitoring reports from the SaaS provider shows that the requests never made it to its systems. The company notices that Amazon CloudWatch Logs are being generated by the Lambda functions. When the same functionality is tested against the EC2 systems, it works as expected.
What is causing the issue?

A. Lambda is in a subnet that does not have a NAT gateway attached to it to connect to the SaaS provider.
B. The end-user application is misconfigured to continue using the endpoint backed by EC2 instances.
C. The throttle limit set on API Gateway is too low and the requests are not making their way through.
D. API Gateway does not have the necessary permissions to invoke Lambda.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 103
A company uses AWS Organizations to manage more than 1,000 AWS accounts. The company has created a new developer organization. There are 540 developer member accounts that must be moved to the new developer organization.
All accounts are set up with all the required information so that each account can be operated as a standalone account.
Which combination of steps should a solutions architect take to move all of the developer accounts to the new developer organization? (Choose three.)

A. Call the MoveAccount operation in the Organizations API from the old organization's management account to migrate the developer accounts to the new developer organization.
B. From the management account, remove each developer account from the old organization using the RemoveAccountFromOrganization operation in the Organizations API.
C. From each developer account, remove the account from the old organization using the RemoveAccountFromOrganization operation in the Organizations API.

D. Sign in to the new developer organization's management account and create a placeholder member account that acts as a target for the developer account migration.
E. Call the InviteAccountToOrganization operation in the Organizations API from the new developer organization's management account to send invitations to the developer accounts.
F. Have each developer sign in to their account and confirm to join the new developer organization.

Correct Answer: CEF
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/organizations/latest/APIReference/API_InviteAccountToOrganization.html

QUESTION 104
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: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 105
A company plans to deploy a new private intranet service on Amazon EC2 instances inside a VPC. An AWS Site-to-Site VPN connects the VPC to the company's on-premises network. The new service must communicate with existing onpremises services. The on-premises services are accessible through the use of hostnames that reside in the company.example DNS zone. This DNS zone is wholly hosted on premises and is available only on the company's private network.
A solutions architect must ensure that the new service can resolve hostnames on the company example domain to integrate with existing services.
Which solution meets these requirements?

A. Create an empty private zone in Amazon Route 53 for company example. Add an additional NS record to the company's on-premises company.example zone that points to the authoritative name servers for the new private zone in Route 53.
B. Turn on DNS hostnames for the VPC. Configure a new outbound endpoint with Amazon Route 53 Resolver. Create a Resolver rule to forward requests for company.example to the on-premises name servers.
C. Turn on DNS hostnames for the VPConfigure a new inbound resolver endpoint with Amazon Route 53 Resolver.
Configure the on-premises DNS server to forward requests for company.example to the new resolver.
D. Use AWS Systems Manager to configure a run document that will install a hosts file that contains any required hostnames. Use an Amazon EventBridge (Amazon CloudWatch Events) rule to run the document when an instance is entering the running state.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Inbound endpoint: DNS resolvers on your network can forward DNS queries to Route 53 Resolver via this endpoint.
Reference: https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver.html

QUESTION 106
A company has a standard three-tier architecture using two Availability Zones. During the company's off season, users report that the website is not working. The Solutions Architect finds that no changes have been made to the environment recently, the website is reachable, and it is possible to log in. However, when the Solutions Architect selects the "find a store near you" function, the maps provided on the site by a third-party RESTful API call do not work about 50% of the time after refreshing the page. The outbound API calls are made through Amazon EC2 NAT instances.
What is the MOST likely reason for this failure and how can it be mitigated in the future?

A. The network ACL for one subnet is blocking outbound web traffic. Open the network ACL and prevent administration from making future changes through IAM.
B. The fault is in the third-party environment. Contact the third party that provides the maps and request a fix that will provide better uptime.
C. One NAT instance has become overloaded. Replace both EC2 NAT instances with a larger-sized instance and make sure to account for growth when making the new instance size.
D. One of the NAT instances failed. Recommend replacing the EC2 NAT instances with a NAT gateway.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The issue is 50% failure, means the balancing over 2 AZs is failing on one NAT instance in one AZ. The solution is to replace the NAT instance with fully managed and high available NAT gateway.

QUESTION 107
A company with multiple accounts is currently using a configuration that does not meet the following security governance policies:
Prevent ingress from port 22 to any Amazon EC2 instance.
Require billing and application tags for resources. Encrypt all Amazon EBS volumes.
A solutions architect wants to provide preventive and detective controls, including notifications about a specific resource, if there are policy deviations.
Which solution should the solutions architect implement?

A. Create an AWS CodeCommit repository containing policy-compliant AWS CloudFormation templates. Create an AWS Service Catalog portfolio. Import the CloudFormation templates by attaching the CodeCommit repository to the portfolio.
Restrict users across all accounts to items from the AWS Service Catalog portfolio. Use AWS Config managed rules to detect deviations from the policies. Configure an Amazon CloudWatch Events rule for deviations, and associate a CloudWatch alarm to send notifications when the TriggeredRules metric is greater than zero.
B. Use AWS Service Catalog to build a portfolio with products that are in compliance with the governance policies in a central account. Restrict users across all accounts to AWS Service Catalog products. Share a compliant portfolio to other accounts. Use AWS Config managed rules to detect deviations from the policies. Configure an Amazon CloudWatch Events rule to send a notification when a deviation occurs.
C. Implement policy-compliant AWS CloudFormation templates for each account, and ensure that all provisioning is completed by CloudFormation. Configure Amazon Inspector to perform regular checks against resources. Perform policy validation and write the assessment output to Amazon CloudWatch Logs. Create a CloudWatch Logs metric filter to increment a metric when a deviation occurs.
Configure a CloudWatch alarm to send notifications when the configured metric is greater than zero.
D. Restrict users and enforce least privilege access using AWS IAM. Consolidate all AWS CloudTrail logs into a single account. Send the CloudTrail logs to Amazon Elasticsearch Service (Amazon ES). Implement monitoring, alerting, and reporting using the Kibana dashboard in Amazon ES and with Amazon SNS.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 108

A company owns a chain of travel agencies and is running an application in the AWS Cloud. Company employees use the application to search for information about travel destinations. Destination content is updated four times each year.
Two fixed Amazon EC2 instances serve the application. The company uses an Amazon Route 53 public hosted zone with a multivalue record of travel.example.com that returns the Elastic IP addresses for the EC2 instances. The application uses Amazon DynamoDB as its primary data store. The company uses a self-hosted Redis instance as a caching solution.
During content updates, the load on the EC2 instances and the caching solution increases drastically. This increased load has led to downtime on several occasions. A solutions architect must update the application so that the application is highly available and can handle the load that is generated by the content updates.
Which solution will meet these requirements?

A. Set up DynamoDB Accelerator (DAX) as in-memory cache. Update the application to use DAX. Create an Auto Scaling group for the EC2 instances. Create an Application Load Balancer (ALB). Set the Auto Scaling group as a target for the ALB.
Update the Route 53 record to use a simple routing policy that targets the ALB's DNS alias. Configure scheduled scaling for the EC2 instances before the content updates.

B. Set up Amazon ElastiCache for Redis. Update the application to use ElastiCache. Create an Auto Scaling group for the EC2 instances. Create an Amazon CloudFront distribution, and set the Auto Scaling group as an origin for the distribution.
Update the Route 53 record to use a simple routing policy that targets the CloudFront distribution's DNS alias. Manually scale up EC2 instances before the content updates.

C. Set up Amazon ElastiCache for Memcached. Update the application to use ElastiCache Create an Auto Scaling group for the EC2 instances. Create an Application Load Balancer (ALB). Set the Auto Scaling group as a target for the ALB. Update the Route 53 record to use a simple routing policy that targets the ALB's DNS alias. Configure scheduled scaling for the application before the content updates.

D. Set up DynamoDB Accelerator (DAX) as in-memory cache. Update the application to use DAX. Create an Auto Scaling group for the EC2 instances. Create an Amazon CloudFront distribution, and set the Auto Scaling group as an origin for the distribution. Update the Route 53 record to use a simple routing policy that targets the CloudFront distribution's DNS alias.
Manually scale up EC2 instances before the content updates.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/dynamodb/dax/

QUESTION 109
Regarding Amazon SNS, you can send notification messages to mobile devices through any of the following supported push notification services, EXCEPT:

A. Microsoft Windows Mobile Messaging (MWMM)
B. Google Cloud Messaging for Android (GCM)
C. Amazon Device Messaging (ADM)

D. Apple Push Notification Service (APNS)

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In Amazon SNS, you have the ability to send notification messages directly to apps on mobile devices. Notification messages sent to a mobile endpoint can appear in the mobile app as message alerts, badge updates, or even sound alerts.
Microsoft Windows Mobile Messaging (MWMM) doesn't exist and is not supported by Amazon SNS.
Reference: http://docs.aws.amazon.com/sns/latest/dg/SNSMobilePush.html

QUESTION 110
A company runs an ordering system on AWS using Amazon SQS and AWS Lambda, with each order received as a JSON message. Recently the company had a marketing event that led to a tenfold increase in orders. With this increase, the following undesired behaviors started in the ordering system:
Lambda failures while processing orders lead to queue backlogs. The same orders have been processed multiple times.
A Solutions Architect has been asked to solve the existing issues with the ordering system and add the following resiliency features:
Retain problematic orders for analysis.
Send notification if errors go beyond a threshold value.
How should the Solutions Architect meet these requirements?

A. Receive multiple messages with each Lambda invocation, add error handling to message processing code and delete messages after processing, increase the visibility timeout for the messages, create a dead letter queue for messages that could not be processed, create an Amazon CloudWatch alarm on Lambda errors for notification.
B. Receive single messages with each Lambda invocation, put additional Lambda workers to poll the queue, delete messages after processing, increase the message timer for the messages, use Amazon CloudWatch Logs for messages that could not be processed, create a CloudWatch alarm on Lambda errors for notification.
C. Receive multiple messages with each Lambda invocation, use long polling when receiving the messages, log the errors from the message processing code using Amazon CloudWatch Logs, create a dead letter queue with AWS Lambda to capture failed invocations, create CloudWatch events on Lambda errors for notification.
D. Receive multiple messages with each Lambda invocation, add error handling to message processing code and delete messages after processing, increase the visibility timeout for the messages, create a delay queue for messages that could not be processed, create an Amazon CloudWatch metric on Lambda errors for notification.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 111
A solutions architect at a large company needs to set up network security for outbound traffic to the internet from all AWS accounts within an organization in AWS Organizations. The organization has more than 100 AWS accounts, and the accounts route to each other by using a centralized AWS Transit Gateway. Each account has both an internet gateway and a NAT gateway for outbound traffic to the internet. The company deploys resources only into a single AWS Region.
The company needs the ability to add centrally managed rule-based filtering on all outbound traffic to the internet for all AWS accounts in the organization. The peak load of outbound traffic will not exceed 25 Gbps in each Availability Zone.
Which solution meets these requirements?

A. Create a new VPC for outbound traffic to the internet. Connect the existing transit gateway to the new VPC. Configure a new NAT gateway. Create an Auto Scaling group of Amazon EC2 instances that run an open-source internet proxy for rulebased filtering across all Availability Zones in the Region. Modify all default routes to point to the proxy's Auto Scaling group.
B. Create a new VPC for outbound traffic to the internet. Connect the existing transit gateway to the new VPC. Configure a new NAT gateway. Use an AWS Network Firewall firewall for rule-based filtering. Create Network Firewall endpoints in each Availability Zone. Modify all default routes to point to the Network Firewall endpoints.
C. Create an AWS Network Firewall firewall for rule-based filtering in each AWS account. Modify all default routes to point to the Network Firewall firewalls in each account.
D. In each AWS account, create an Auto Scaling group of network-optimized Amazon EC2 instances that run an opensource internet proxy for rule-based filtering. Modify all default routes to point to the proxy's Auto Scaling group.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 112
Which is a valid Amazon Resource name (ARN) for IAM?

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 113
You have setup an Auto Scaling group. The cool down period for the Auto Scaling group is 7 minutes. The first scaling activity request for the Auto Scaling group is to launch two instances. It receives the activity question at time "t", and the first instance is launched at t+3 minutes, while the second instance is launched at t +4 minutes.
How many minutes after time "t" will Auto Scaling accept another scaling activity request?

A. 11 minutes
B. 10 minutes
C. 7 minutes
D. 14 minutes

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
If an Auto Scaling group is launching more than one instance, the cool down period for each instance starts after that instance is launched. The group remains locked until the last instance that was launched has completed its cool down period. In this case the cool down period for the first instance starts after 3 minutes and finishes at the 10th minute (3+7 cool down), while for the second instance it starts at the 4th minute and finishes at the 11th minute (4+7 cool down). Thus, the Auto Scaling group will receive another request only after 11 minutes.
Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AS_Concepts.html

QUESTION 114
A media company is serving video files stored in Amazon S3 using Amazon CloudFront. The development team needs access to the logs to diagnose faults and perform service monitoring. The log files from CloudFront may contain sensitive information about users.
The company uses a log processing service to remove sensitive information before making the logs available to the development team. The company has the following requirements for the unprocessed logs:

The logs must be encrypted at rest and must be accessible by the log processing service only.
Only the data protection team can control access to the unprocessed log files.
AWS CloudFormation templates must be stored in AWS CodeCommit.
AWS CodePipeline must be triggered on commit to perform updates made to CloudFormation templates.
CloudFront is already writing the unprocessed logs to an Amazon S3 bucket, and the log processing service is operating against this S3 bucket.
Which combination of steps should a solutions architect take to meet the company's requirements? (Choose two.)

A. Create an AWS KMS key that allows the AWS Logs Delivery account to generate data keys for encryption Configure S3 default encryption to use server-side encryption with KMS managed keys (SSE-KMS) on the log storage bucket using the new KMS key. Modify the KMS key policy to allow the log processing service to perform decrypt operations.
B. Create an AWS KMS key that follows the CloudFront service role to generate data keys for encryption Configure S3 default encryption to use KMS managed keys (SSE-KMS) on the log storage bucket using the new KMS key Modify the KMS key policy to allow the log processing service to perform decrypt operations.
C. Configure S3 default encryption to use AWS KMS managed keys (SSE-KMS) on the log storage bucket using the AWS Managed S3 KMS key. Modify the KMS key policy to allow the CloudFront service role to generate data keys for encryption Modify the KMS key policy to allow the log processing service to perform decrypt operations.
D. Create a new CodeCommit repository for the AWS KMS key template.
Create an IAM policy to allow commits to the new repository and attach it to the data protection team's users. Create a new CodePipeline pipeline with a custom IAM role to perform KMS key updates using CloudFormation Modify the KMS key policy to allow the CodePipeline IAM role to modify the key policy.
E. Use the existing CodeCommit repository for the AWS KMS key template.
Create an IAM policy to allow commits to the new repository and attach it to the data protection team's users.
Modify the existing CodePipeline pipeline to use a custom IAM role and to perform KMS key updates using CloudFormation.
Modify the KMS key policy to allow the CodePipeline IAM role to modify the key policy.

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 115
A web application is hosted in a dedicated VPC that is connected to a company's on-premises data center over a Site-to-Site VPN connection. The application is accessible from the company network only. This is a temporary non-production application that is used during business hours. The workload is generally low with occasional surges.
The application has an Amazon Aurora MySQL provisioned database cluster on the backend. The VPC has an internet gateway and a NAT gateways attached. The web servers are in private subnets in an Auto Scaling group behind an Elastic Load Balancer. The web servers also upload data to an Amazon S3 bucket through the internet.
A solutions architect needs to reduce operational costs and simplify the architecture.
Which strategy should the solutions architect use?

A. Review the Auto Scaling group settings and ensure the scheduled actions are specified to operate the Amazon EC2 instances during business hours only. Use 3-year scheduled Reserved Instances for the web server EC2 instances. Detach the internet gateway and remove the NAT gateways from the VPC. Use an Aurora Serverless database and set up a VPC endpoint for the S3 bucket.

B. Review the Auto Scaling group settings and ensure the scheduled actions are specified to operate the Amazon EC2 instances during business hours only. Detach the internet gateway and remove the NAT gateways from the VPC. Use an Aurora Serverless database and set up a VPC endpoint for the S3 bucket, then update the network routing and security rules and policies related to the changes.

C. Review the Auto Scaling group settings and ensure the scheduled actions are specified to operate the Amazon EC2 instances during business hours only. Detach the internet gateway from the VPC, and use an Aurora Serverless database.
Set up a VPC endpoint for the S3 bucket, then update the network routing and security rules and policies related to the changes.

D. Use 3-year scheduled Reserved Instances for the web server Amazon EC2 instances. Remove the NAT gateways from the VPC, and set up a VPC endpoint for the S3 bucket. Use Amazon CloudWatch and AWS Lambda to stop and start the Aurora DB cluster so it operates during business hours only. Update the network routing and security rules and policies related to the changes.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 116
A company is migrating applications from on premises to the AWS Cloud. These applications power the company's internal web forms. These web forms collect data for specific events several times each quarter. The web forms use simple SQL statements to save the data to a local relational database.
Data collection occurs for each event, and the on-premises servers are idle most of the time. The company needs to minimize the amount of idle infrastructure that supports the web forms.
Which solution will meet these requirements?

A. Use Amazon EC2 Image Builder to create AMIs for the legacy servers. Use the AMIs to provision EC2 instances to recreate the applications in the AWS Cloud. Place an Application Load Balancer (ALB) in front of the EC2 instances. Use Amazon Route 53 to point the DNS names of the web forms to the ALB.

B. Create one Amazon DynamoDB table to store data for all the data input. Use the application form name as the table key to distinguish data items. Create an Amazon Kinesis data stream to receive the data input and store the input in DynamoDUse Amazon Route 53 to point the DNS names of the web forms to the Kinesis data stream's endpoint.

C. Create Docker images for each server of the legacy web form applications. Create an Amazon Elastic Container Service (Amazon EC2) cluster on AWS Fargate. Place an Application Load Balancer in front of the ECS cluster. Use Fargate task storage to store the web form data.

D. Provision an Amazon Aurora Serverless cluster. Build multiple schemas for each web form's data storage. Use Amazon API Gateway and an AWS Lambda function to recreate the data input forms. Use Amazon Route 53 to point the DNS names of the web forms to their corresponding API Gateway endpoint.

Correct Answer: B

Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/kds.html

QUESTION 117
A company has several AWS accounts. A development team is building an automation framework for cloud governance and remediation processes. The automation framework uses AWS Lambda functions in a centralized account. A solutions architect must implement a least privilege permissions policy that allows the Lambda functions to run in each of the company's AWS accounts.
Which combination of steps will meet these requirements? (Choose two.)

A. In the centralized account, create an IAM role that has the Lambda service as a trusted entity. Add an inline policy to assume the roles of the other AWS accounts.
B. In the other AWS accounts, create an IAM role that has minimal permissions. Add the centralized account's Lambda IAM role as a trusted entity.
C. In the centralized account, create an IAM role that has roles of the other accounts as trusted entities. Provide minimal permissions.
D. In the other AWS accounts, create an IAM role that has permissions to assume the role of the centralized account. Add the Lambda service as a trusted entity.
E. In the other AWS accounts, create an IAM role that has minimal permissions. Add the Lambda service as a trusted entity.

Correct Answer: AC
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/blogs/devops/how-to-centrally-manage-aws-config-rules-across-multiple-aws-accounts/

QUESTION 118
An organization is setting up a highly scalable application using Elastic Beanstalk. The organization is using ELB and RDS with VPC. The organization has public and private subnets within the cloud.
Which of the below mentioned configurations will not work in this scenario?

A. To setup RDS in a private subnet and ELB in a public subnet.
B. The configuration must have public and private subnets in the same AZ.
C. The configuration must have two private subnets in separate AZs.
D. The EC2 instance should have a public IP assigned to it.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. If the organization is planning to implement a scalable secure application using RDS, VPC and ELB the organization should follow below mentioned configurations:
Setup RDS in a private subnet Setup ELB in a public subnet
Since RDS needs a subnet group, the organization should have two private subnets in the same zone The ELB needs private and public subnet to be part of same AZs It is not required that instances should have a public IP assigned to them. The instances can be a part of a private subnet and the organization can setup a corresponding routing mechanism.
Reference: http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/vpc-rds.html

QUESTION 119
A company recently transformed its legacy infrastructure provisioning scripts to AWS CloudFormation templates. The newly developed templates are hosted in the company's private GitHub repository. Since adopting CloudFormation, the company has encountered several issues with updates to the CloudFormation templates, causing execution or creating environment.
Management is concerned by the increase in errors and has asked a Solutions Architect to design the automated testing of CloudFormation template updates.
What should the Solution Architect do to meet these requirements?

A. Use AWS CodePipeline to create a change set from the CloudFormation templates stored in the private GitHub repository. Execute the change set using AWS CodeDeploy. Include a CodePipeline action to test the deployment with testing scripts run by AWS CodeBuild.
B. Mirror the GitHub repository to AWS CodeCommit using AWS Lambda. Use AWS CodeDeploy to create a change set from the CloudFormation templates and execute it. Have CodeDeploy test the deployment with testing scripts run by AWS CodeBuild.
C. Use AWS CodePipeline to create and execute a change set from the CloudFormation templates stored in the GitHub repository. Configure a CodePipeline action to be deployment with testing scripts run by AWS CodeBuild.
D. Mirror the GitHub repository to AWS CodeCommit using AWS Lambda. Use AWS CodeBuild to create a change set from the CloudFormation templates and execute it. Have CodeBuild test the deployment with testing scripts.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 120
A company has developed a mobile game. The backend for the game runs on several virtual machines located in an onpremises data center. The business logic

is exposed using a REST API with multiple functions. Player session data is stored in central file storage. Backend services use different API keys for throttling and to distinguish between live and test traffic.
The load on the game backend varies throughout the day. During peak hours, the server capacity is not sufficient. There are also latency issues when fetching player session data. Management has asked a solutions architect to present a cloud architecture that can handle the game's varying load and provide low-latency data access. The API model should not be changed.
Which solution meets these requirements?

A. Implement the REST API using a Network Load Balancer (NLB). Run the business logic on an Amazon EC2 instance behind the NLB. Store player session data in Amazon Aurora Serverless.
B. Implement the REST API using an Application Load Balancer (ALB). Run the business logic in AWS Lambda. Store player session data in Amazon DynamoDB with on-demand capacity.
C. Implement the REST API using Amazon API Gateway. Run the business logic in AWS Lambda. Store player session data in Amazon DynamoDB with on-demand capacity.
D. Implement the REST API using AWS AppSync. Run the business logic in AWS Lambda. Store player session data in Amazon Aurora Serverless.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 121
A company has built a high performance computing (HPC) cluster in AWS for a tightly coupled workload that generates a large number of shared files stored in Amazon EFS. The cluster was performing well when the number of Amazon EC2 instances in the cluster was 100. However, when the company increased the cluster size to 1,000 EC2 instances, overall performance was well below expectations. Which collection of design choices should a solutions architect make to achieve the maximum performance from the HPC cluster? (Choose three.)

A. Ensure the HPC cluster is launched within a single Availability Zone.
B. Launch the EC2 instances and attach elastic network interfaces in multiples of four.
C. Select EC2 instance types with an Elastic Fabric Adapter (EFA) enabled.
D. Ensure the clusters is launched across multiple Availability Zones.
E. Replace Amazon EFS win multiple Amazon EBS volumes in a RAID array.
F. Replace Amazon EFS with Amazon FSx for Lustre.

Correct Answer: DEF
Section: (none)
Explanation

Explanation/Reference:

QUESTION 122
A Solutions Architect is building a solution for updating user metadata that is initiated by web servers. The solution needs to rapidly scale from hundreds to tens of thousands of jobs in less than 30 seconds. The solution must be asynchronous always avertable and minimize costs.
Which strategies should the Solutions Architect use to meet these requirements?

A. Create an AWS SWF worker that will update user metadata updating web application to start a new workflow for every job.
B. Create an AWS Lambda function that will update user metadata. Create an Amazon SOS queue and configure it as an event source for the Lambda function. Update the web application to send jobs to the queue.
C. Create an AWS Lambda function that will update user metadata. Create AWS Step Functions that will trigger the Lambda function. Update the web application to initiate Step Functions for every job.
D. Create an Amazon SQS queue. Create an AMI with a worker to check the queue and update user metadata. Configure an Amazon EC2 Auto Scaling group with the new AMI. Update the web application to send jobs to the queue.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 123
A company's security compliance requirements state that all Amazon EC2 images must be scanned for vulnerabilities and must pass a CVE assessment. A solutions architect is developing a mechanism to create security- approved AMIs that can be used by developers. Any new AMIs should go through an automated assessment process and be marked as approved before developers can use them. The approved images must be scanned every 30 days to ensure compliance.
Which combination of steps should the solutions architect take to meet these requirements while following best practices?
(Choose two.)

A. Use the AWS Systems Manager EC2 agent to run the CVE assessment on the EC2 instances launched from the AMIs that need to be scanned.
B. Use AWS Lambda to write automatic approval rules. Store the approved AMI list in AWS Systems Manager Parameter Store. Use Amazon EventBridge to trigger an AWS Systems Manager Automation document on all EC2 instances every 30 days.
C. Use Amazon Inspector to run the CVE assessment on the EC2 instances launched from the AMIs that need to be scanned.
D. Use AWS Lambda to write automatic approval rules. Store the approved AMI list in AWS Systems Manager Parameter Store. Use a managed AWS Config rule for continuous scanning on all EC2 instances, and use AWS Systems Manager Automation documents for remediation.
E. Use AWS CloudTrail to run the CVE assessment on the EC2 instances launched from the AMIs that need to be scanned.

Correct Answer: BC
Section: (none)
Explanation

Explanation/Reference:

QUESTION 124
An organization is setting up a highly scalable application using Elastic Beanstalk.
They are using Elastic Load Balancing (ELB) as well as a Virtual Private Cloud (VPC) with public and private subnets. They have the following requirements:
- All the EC2 instances should have a private IP
- All the EC2 instances should receive data via the ELB's.
Which of these will not be needed in this setup?

A. Launch the EC2 instances with only the public subnet.
B. Create routing rules which will route all inbound traffic from ELB to the EC2 instances.
C. Configure ELB and NAT as a part of the public subnet only.
D. Create routing rules which will route all outbound traffic from the EC2 instances through NAT.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. If the organization wants the Amazon EC2 instances to have a private IP address, he should create a public and private subnet for VPC in each Availability Zone (this is an AWS Elastic Beanstalk requirement). The organization should add their public resources, such as ELB and NAT to the public subnet, and AWC Elastic Beanstalk will assign them unique elastic IP addresses (a static, public IP address). The organization should launch Amazon EC2 instances in a private subnet so that AWS Elastic Beanstalk assigns them non-routable private IP addresses. Now the organization should configure route tables with the following rules: route all inbound traffic from ELB to EC2 instances route all outbound traffic from EC2 instances through NAT
Reference: http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/AWSHowTo-vpc.html

QUESTION 125
The CFO of a company wants to allow one of his employees to view only the AWS usage report page.
Which of the below mentioned IAM policy statements allows the user to have access to the AWS usage report page?

A. "Effect": "Allow", "Action": ["Describe"], "Resource": "Billing"
B. "Effect": "Allow", "Action": ["aws-portal: ViewBilling"], "Resource": "*"

C. "Effect": "Allow", "Action": ["aws-portal: ViewUsage"], "Resource": "*"
D. "Effect": "Allow", "Action": ["AccountUsage], "Resource": "*"

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If the CFO wants to allow only AWS usage report page access, the policy for that IAM user will be as given below:
{
"Version": "2012-10-17",
"Statement": [
{
"Effect": "Allow", "Action": [
"aws-portal:ViewUsage"
],
"Resource": "*"
}
]
}
Reference: http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/billing-permissions-ref.html

QUESTION 126
A company runs an application in the cloud that consists of a database and a website. Users can post data to the website, have the data processed, and have the data sent back to them in an email. Data is stored in a MySQL database running on an Amazon EC2 instance. The database is running in a VPC with two private subnets. The website is running on Apache Tomcat in a single EC2 instance in a different VPC with one public subnet. There is a single VPC peering connection between the database and website VPC.
The website has suffered several outages during the last month due to high traffic.
Which actions should a solutions architect take to increase the reliability of the application? (Choose three.)

A. Place the Tomcat server in an Auto Scaling group with multiple EC2 instances behind an Application Load Balancer.
B. Provision an additional VPC peering connection.
C. Migrate the MySQL database to Amazon Aurora with one Aurora Replica.
D. Provision two NAT gateways in the database VPC.
E. Move the Tomcat server to the database VPC.
F. Create an additional public subnet in a different Availability Zone in the website VPC.

Correct Answer: ACF

Section: (none)
Explanation

Explanation/Reference:

QUESTION 127
A company runs an application that gives users the ability to search for videos and related information by using keywords that are curated from content providers. The application data is stored in an onpremises Oracle database that is 800 GB in size.
The company wants to migrate the data to an Amazon Aurora MySQL DB instance. A solutions architect plans to use the AWS Schema Conversion Tool and AWS Database Migration Service (AWS DMS) for the migration. During the migration, the existing database must serve ongoing requests. The migration must be completed with minimum downtime.
Which solution will meet these requirements?

A. Create primary key indexes, secondary indexes, and referential integrity constraints in the target database before starting the migration process.
B. Use AWS DMS to run the conversion report for Oracle to Aurora MySQL. Remediate any issues. Then use AWS DMS to migrate the data.
C. Use the M5 or C5 DMS replication instance type for ongoing replication.
D. Turn off automatic backups and logging of the target database until the migration and cutover processes are complete.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/dms/latest/sbs/chap-rdsoracle2aurora.html

QUESTION 128
One of your AWS Data Pipeline activities has failed consequently and has entered a hard failure state after retrying thrice.
You want to try it again. Is it possible to increase the number of automatic retries to more than thrice?

A. Yes, you can increase the number of automatic retries to 6.
B. Yes, you can increase the number of automatic retries to indefinite number.
C. No, you cannot increase the number of automatic retries.
D. Yes, you can increase the number of automatic retries to 10.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In AWS Data Pipeline, an activity fails if all of its activity attempts return with a failed state. By default, an activity retries three times before entering a hard failure state. You can increase the number of automatic retries to 10. However, the system does not allow indefinite retries.
Reference:
https://aws.amazon.com/datapipeline/faqs/

QUESTION 129
Your Application is not highly available, and your on-premises server cannot access the mount target because the Availability Zone (AZ) in which the mount target exists is unavailable.
Which of the following actions is recommended?

A. The application must implement the checkpoint logic and recreate the mount target.
B. The application must implement the shutdown logic and delete the mount target in the AZ.
C. The application must implement the delete logic and connect to a different mount target in the same AZ.
D. The application must implement the restart logic and connect to a mount target in a different AZ.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
To make sure that there is continuous availability between your on-premises data center and your Amazon Virtual Private Cloud (VPC), it is suggested that you configure two AWS Direct Connect connections. Your application should implement restart logic and connect to a mount target in a different AZ if your application is not highly available and your on-premises server cannot access the mount target because the AZ in which the mount target exists becomes unavailable.
Reference: http://docs.aws.amazon.com/efs/latest/ug/performance.html#performance-onpremises

QUESTION 130
An organization has created multiple components of a single application for compartmentalization. Currently all the components are hosted on a single EC2 instance. Due to security reasons the organization wants to implement two separate SSLs for the separate modules although it is already using VPC.
How can the organization achieve this with a single instance?

A. You have to launch two instances each in a separate subnet and allow VPC peering for a single IP.
B. Create a VPC instance which will have multiple network interfaces with multiple elastic IP addresses.
C. Create a VPC instance which will have both the ACL and the security group attached to it and have separate rules for each IP address.
D. Create a VPC instance which will have multiple subnets attached to it and each will have a separate IP address.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. With VPC the user can specify multiple private IP addresses for his instances.
The number of network interfaces and private IP addresses that a user can specify for an instance depends on the instance type. With each network interface the organization can assign an EIP. This scenario helps when the user wants to host multiple websites on a single EC2 instance by using multiple SSL certificates on a single server and associating each certificate with a specific EIP address. It also helps in scenarios for operating network appliances, such as firewalls or load balancers that have multiple private IP addresses for each network interface.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/MultipleIP.html

QUESTION 131
A solutions architect is implementing infrastructure as code for a two-tier web application in an AWS CloudFormation template. The web frontend application will be deployed on Amazon EC2 instances in an Auto Scaling group. The backend database will be an Amazon RDS for MySQL DB instance. The database password will be rotated every 60 days.
How can the solutions architect MOST securely manage the configuration of the application's database credentials?

A. Provide the database password as a parameter in the CloudFormation template. Create an initialization script in the Auto Scaling group's launch configuration UserData property to reference the password parameter using the Ref intrinsic function.
Store the password on the EC2 instances. Reference the parameter for the value of the MasterUserPassword property in the AWS::RDS::DBInstance resource using the Ref intrinsic function.
B. Create a new AWS Secrets Manager secret resource in the CloudFormation template to be used as the database password. Configure the application to retrieve the password from Secrets Manager when needed. Reference the secret resource for the value of the MasterUserPassword property in the AWS::RDS::DBInstance resource using a dynamic reference.
C. Create a new AWS Secrets Manager secret resource in the CloudFormation template to be used as the database password. Create an initialization script in the Auto Scaling group's launch configuration UserData property to reference the secret resource using the Ref intrinsic function. Reference the secret resource for the value of the MasterUserPassword property in the AWS::RDS::DBInstance resource using the Ref intrinsic function.
D. Create a new AWS Systems Manager Parameter Store parameter in the CloudFormation template to be used as the database password. Create an initialization script in the Auto Scaling group's launch configuration UserData property to reference the parameter. Reference the parameter for the value of the MasterUserPassword property in the AWS::RDS::DBInstance resource using the Fn::GetAtt intrinsic function.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 132
A company is developing a web application that runs on Amazon EC2 instances in an Auto Scaling group behind a publicfacing Application Load Balancer (ALB). Only users from a specific country are allowed to access the application. The company needs the ability to log the access requests that have been blocked. The solution should require the least possible maintenance.
Which solution meets these requirements?

A. Create an IPSet containing a list of IP ranges that belong to the specified country. Create an AWS WAF web ACL. Configure a rule to block any requests that do not originate from an IP range in the IPSet. Associate the rule with the web ACL. Associate the web ACL with the ALB.
B. Create an AWS WAF web ACL. Configure a rule to block any requests that do not originate from the specified country. Associate the rule with the web ACL. Associate the web ACL with the ALB.
C. Configure AWS Shield to block any requests that do not originate from the specified country. Associate AWS Shield with the ALB.
D. Create a security group rule that allows ports 80 and 443 from IP ranges that belong to the specified country. Associate the security group with the ALB.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 133
A user is creating a Provisioned IOPS volume. What is the maximum ratio the user should configure between Provisioned IOPS and the volume size?

A. 30 to 1
B. 50 to 1
C. 10 to 1
D. 20 to 1

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Provisioned IOPS SSD (io1) volumes are designed to meet the needs of I/O-intensive workloads, particularly database workloads, that are sensitive to storage performance and consistency. An io1 volume can range in size from 4 GiB to 16 TiB and you can provision 100 up to 20,000 IOPS per volume. The maximum

ratio of provisioned IOPS to requested volume size (in GiB) is 50:1. For example, a 100 GiB volume can be provisioned with up to 5,000 IOPS. Any volume 400 GiB in size or greater allows provisioning up to the 20,000 IOPS maximum.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html

QUESTION 134
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.
Which approach can satisfy these objectives?

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 135
A user has suspended the scaling process on the Auto Scaling group. A scaling activity to increase the instance count was already in progress.
What effect will the suspension have on that activity?

A. No effect. The scaling activity continues
B. Pauses the instance launch and launches it only after Auto Scaling is resumed
C. Terminates the instance
D. Stops the instance temporary

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The user may want to stop the automated scaling processes on the Auto Scaling groups either to perform manual operations or during emergency situations. To

perform this, the user can suspend one or more scaling processes at any time. When this process is suspended, Auto Scaling creates no new scaling activities for that group. Scaling activities that were already in progress before the group was suspended continue until completed.
Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AS_Concepts.html

QUESTION 136
An ecommerce company has an order processing application it wants to migrate to AWS. The application has inconsistent data volume patterns, but needs to be avail at all times. Orders must be processed as they occur and in the order that they are received.
Which set of steps should a solutions architect take to meet these requirements?

A. Use AWS Transfer for SFTP and upload orders as they occur. Use On-Demand Instances in multiple Availability Zones for processing.
B. Use Amazon SNS with FIFO and send orders as they occur. Use a single large Reserved Instance for processing.
C. Use Amazon SQS with FIFO and send orders as they occur. Use Reserved Instances in multiple Availability Zones for processing.
D. Use Amazon SQS with FIFO and send orders as they occur. Use Spot Instances in multiple Availability Zones for processing.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 137
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:

Explanation:
Reference: http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/AccessLogs.html

QUESTION 138
A media storage application uploads user photos to Amazon S3 for processing. End users are reporting that some uploaded photos are not being processed properly. The Application Developers trace the logs and find that AWS Lambda is experiencing execution issues when thousands of users are on the system simultaneously. Issues are caused by:
Limits around concurrent executions.
The performance of Amazon DynamoDB when saving data.
Which actions can be taken to increase the performance and reliability of the application? (Choose two.)

A. Evaluate and adjust the read capacity units (RCUs) for the DynamoDB tables.
B. Evaluate and adjust the write capacity units (WCUs) for the DynamoDB tables.
C. Add an Amazon ElastiCache layer to increase the performance of Lambda functions.
D. Configure a dead letter queue that will reprocess failed or timed-out Lambda functions.
E. Use S3 Transfer Acceleration to provide lower-latency access to end users.

Correct Answer: BD
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
B: https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.htmlD:
https://aws.amazon.com/blogs/compute/robust-serverless-application-design-with-aws-lambda-dlq/

QUESTION 139
A company is creating a centralized logging service running on Amazon EC2 that will receive and analyze logs from hundreds of AWS accounts. AWS PrivateLink is being used to provide connectivity between the client services and the logging service.
In each AWS account with a client an interface endpoint has been created for the logging service and is available. The logging service running on EC2 instances with a Network Load Balancer (NLB) are deployed in different subnets. The clients are unable to submit logs using the VPC endpoint.
Which combination of steps should a solutions architect take to resolve this issue? (Choose two.)

A. Check that the NACL is attached to the logging service subnet to allow communications to and from the NLB subnets.
Check that the NACL is attached to the NLB subnet to allow communications to and from the logging service subnets running on EC2 instances.
B. Check that the NACL is attached to the logging service subnets to allow communications to and from the interface endpoint subnets. Check that the NACL is attached to the interface endpoint subnet to allow communications to and from the logging service subnets running on EC2 instances.
C. Check the security group for the logging service running on the EC2 instances to ensure it allows ingress from the NLB subnets.

D. Check the security group for the logging service running on the EC2 instances to ensure it allows ingress from the clients.
E. Check the security group for the NLB to ensure it allows ingress from the interface endpoint subnets.

Correct Answer: DE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 140
A company has asked a Solutions Architect to design a secure content management solution that can be accessed by API calls by external customer applications. The company requires that a customer administrator must be able to submit an API call and roll back changes to existing files sent to the content management solution, as needed.
What is the MOST secure deployment design that meets all solution requirements?

A. Use Amazon S3 for object storage with versioning and bucket access logging enabled, and an IAM role and access policy for each customer application. Encrypt objects using SSE-KMS. Develop the content management application to use a separate AWS KMS key for each customer.
B. Use Amazon WorkDocs for object storage. Leverage WorkDocs encryption, user access management, and version control. Use AWS CloudTrail to log all SDK actions and create reports of hourly access by using the Amazon CloudWatch dashboard. Enable a revert function in the SDK based on a static Amazon S3 webpage that shows the output of the CloudWatch dashboard.
C. Use Amazon EFS for object storage, using encryption at rest for the Amazon EFS volume and a customer managed key stored in AWS KMS. Use IAM roles and Amazon EFS access policies to specify separate encryption keys for each customer application. Deploy the content management application to store all new versions as new files in Amazon EFS and use a control API to revert a specific file to a previous version.
D. Use Amazon S3 for object storage with versioning and enable S3 bucket access logging. Use an IAM role and access policy for each customer application. Encrypt objects using client-side encryption, and distribute an encryption key to all customers when accessing the content management application.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 141
A user is planning to use EBS for his DB requirement. The user already has an EC2 instance running in the VPC private subnet. How can the user attach the EBS volume to a running instance?

A. The user can create EBS in the same zone as the subnet of instance and attach that EBS to instance.

B. It is not possible to attach an EBS to an instance running in VPC until the instance is stopped.
C. The user can specify the same subnet while creating EBS and then attach it to a running instance.
D. The user must create EBS within the same VPC and then attach it to a running instance.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. The VPC is always specific to a region. The user can create a VPC which can span multiple Availability Zones by adding one or more subnets in each Availability Zone. The instance launched will always be in the same availability zone of the respective subnet. When creating an EBS the user cannot specify the subnet or VPC. However, the user must create the EBS in the same zone as the instance so that it can attach the EBS volume to the running instance.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html#VPCSubnet

QUESTION 142
A company has an application that runs a web service on Amazon EC2 instances and stores .jpg images in Amazon S3. The web traffic has a predictable baseline, but often demand spikes unpredictably for short periods of time. The application is loosely coupled and stateless. The .jpg images stored in Amazon S3 are accessed frequently for the first 15 to 20 days, they are seldom accessed thereafter but always need to be immediately available. The CIO has asked to find ways to reduce costs.
Which of the following options will reduce costs? (Choose two.)

A. Purchase Reserved instances for baseline capacity requirements and use On-Demand instances for the demand spikes.
B. Configure a lifecycle policy to move the .jpg images on Amazon S3 to S3 IA after 30 days.
C. Use On-Demand instances for baseline capacity requirements and use Spot Fleet instances for the demand spikes.
D. Configure a lifecycle policy to move the .jpg images on Amazon S3 to Amazon Glacier after 30 days.
E. Create a script that checks the load on all web servers and terminates unnecessary On-Demand instances.

Correct Answer: AB
Section: (none)
Explanation

Explanation/Reference:

QUESTION 143
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 DynamoDB 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 144
A company has more than 10,000 sensors that send data to an on-premises Apache Kafka server by using the Message Queuing Telemetry Transport (MQTT) protocol. The on-premises Kafka server transforms the data and then stores the results as objects in an Amazon S3 bucket.
Recently, the Kafka server crashed. The company lost sensor data while the server was being restored. A solutions architect must create a new design on AWS that is highly available and scalable to prevent a similar occurrence.
Which solution will meet these requirements?

A. Launch two Amazon EC2 instances to host the Kafka server in an active/standby configuration across two Availability Zones. Create a domain name in Amazon Route 53. Create a Route 53 failover policy. Route the sensors to send the data to the domain name.
B. Migrate the on-premises Kafka server to Amazon Managed Streaming for Apache Kafka (Amazon MSK). Create a Network Load Balancer (NLB) that points to the Amazon MSK broker. Enable NLB health checks. Route the sensors to send the data to the NLB.
C. Deploy AWS IoT Core, and connect it to an Amazon Kinesis Data Firehose delivery stream. Use an AWS Lambda function to handle data transformation. Route the sensors to send the data to AWS IoT Core.
D. Deploy AWS IoT Core, and launch an Amazon EC2 instance to host the Kafka server. Configure AWS IoT Core to send the data to the EC2 instance. Route the sensors to send the data to AWS IoT Core.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/msk/faqs/

QUESTION 145
A government client needs you to set up secure cryptographic key storage for some of their extremely confidential data. You decide that the AWS CloudHSM is the best service for this. However, there seem to be a few pre-requisites before this can happen, one of those being a security group that has certain ports open. Which of the following is correct in regards to those security groups?

A. A security group that has no ports open to your network.
B. A security group that has only port 3389 (for RDP) open to your network.
C. A security group that has only port 22 (for SSH) open to your network.
D. A security group that has port 22 (for SSH) or port 3389 (for RDP) open to your network.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
AWS CloudHSM provides secure cryptographic key storage to customers by making hardware security modules (HSMs) available in the AWS cloud.
AWS CloudHSM requires the following environment before an HSM appliance can be provisioned. A virtual private cloud (VPC) in the region where you want the AWS CloudHSM service. One private subnet (a subnet with no Internet gateway) in the VPC. The HSM appliance is provisioned into this subnet.
One public subnet (a subnet with an Internet gateway attached). The control instances are attached to this subnet.
An AWS Identity and Access Management (IAM) role that delegates access to your AWS resources to AWS CloudHSM.
An EC2 instance, in the same VPC as the HSM appliance, that has the SafeNet client software installed. This instance is referred to as the control instance and is used to connect to and manage the HSM appliance.
A security group that has port 22 (for SSH) or port 3389 (for RDP) open to your network. This security group is attached to your control instances so you can access them remotely.

QUESTION 146
To serve Web traffic for a popular product your chief financial officer and IT director have purchased 10 m1.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 m1.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 m1.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 m1.large instances.

C. Route traffic to EC2 m1.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: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html

QUESTION 147
A Solutions Architect needs to migrate a legacy application from on premises to AWS. On premises, the application runs on two Linux servers behind a load balancer and accesses a database that is master-master on two servers. Each application server requires a license file that is tied to the MAC address of the server's network adapter. It takes the software vendor 12 hours to send ne license files through email. The application requires configuration files to use static. IPv4 addresses to access the database servers, not DNS.
Given these requirements, which steps should be taken together to enable a scalable architecture for the application servers? (Choose two.)

A. Create a pool of ENIs, request license files from the vendor for the pool, and store the license files within Amazon S3. Create automation to download an unused license, and attach the corresponding ENI at boot time.

B. Create a pool of ENIs, request license files from the vendor for the pool, store the license files on an Amazon EC2 instance, modify the configuration files, and create an AMI from the instance. use this AMI for all instances.

C. Create a bootstrap automation to request a new license file from the vendor with a unique return email. Have the server configure itself with the received license file.

D. Create bootstrap automation to attach an ENI from the pool, read the database IP addresses from AWS Systems Manager Parameter Store, and inject those parameters into the local configuration files. Keep SSM up to date using a Lambda function.

E. Install the application on an EC2 instance, configure the application, and configure the IP address information. Create an AMI from this instance and use if for all instances.

Correct Answer: CD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 148
What is the maximum length for an instance profile name in AWS IAM?

A. 512 characters
B. 128 characters
C. 1024 characters
D. 64 characters

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The maximum length for an instance profile name is 128 characters.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/LimitationsOnEntities.html

QUESTION 149
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:
Explanation:
A default security group is named default, and it has an ID assigned by AWS. The following are the initial settings for each default security group:
Allow inbound traffic only from other instances associated with the default security group Allow all outbound traffic from the instance The default security group specifies itself as a source security group in its inbound rules. This is what allows instances associated with the default security group to communicate with other instances associated with the default security group.
Reference:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-security-groups.html#default-%20security-group

QUESTION 150

A hybrid network architecture must be used during a company's multi-year data center migration from multiple private data centers to AWS. The current data centers are linked together with private fiber. Due to unique legacy applications, NAT cannot be used. During the migration period, many applications will need access to other applications in both the data centers and AWS.
Which option offers a hybrid network architecture that is secure and highly available, that allows for high bandwidth and a multi-region deployment post-migration?

A. Use AWS Direct Connect to each data center from different ISPs, and configure routing to failover to the other data center's Direct Connect if one fails. Ensure that no VPC CIDR blocks overlap one another or the on-premises network.
B. Use multiple hardware VPN connections to AWS from the on-premises data center. Route different subnet traffic through different VPN connections. Ensure that no VPC CIDR blocks overlap one another or the on-premises network.
C. Use a software VPN with clustering both in AWS and the on-premises data center, and route traffic through the cluster. Ensure that no VPC CIDR blocks overlap one another or the on-premises network.
D. Use AWS Direct Connect and a VPN as backup, and configure both to use the same virtual private gateway and BGP. Ensure that no VPC CIDR blocks overlap one another or the on-premises network.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 151
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 152
A company is building an AWS landing zone and has asked a Solutions Architect to design a multi-account access strategy that will allow hundreds of users to use corporate credentials to access the AWS Console. The company is running a Microsoft Active Directory, and users will use an AWS Direct Connect connection to connect to AWS. The company also wants to be able to federate to third-party services and providers, including custom applications.
Which solution meets the requirements by using the LEAST amount of management overhead?

A. Connect the Active Directory to AWS by using single sign-on and an Active Directory Federation Services (AD FS) with SAML 2.0, and then configure the Identity Provider (IdP) system to use formbased authentication. Build the AD FS portal page with corporate branding, and integrate third-party applications that support SAML 2.0 as required.
B. Create a two-way Forest trust relationship between the on-premises Active Directory and the AWS Directory Service. Set up AWS Single Sign-On with AWS Organizations. Use single sign-on integrations for connections with third-party applications.
C. Configure single sign-on by connecting the on-premises Active Directory using the AWS Directory Service AD Connector.
Enable federation to the AWS services and accounts by using the IAM applications and services linking function. Leverage third-party single sign-on as needed.
D. Connect the company's Active Directory to AWS by using AD FS and SAML 2.0. Configure the AD FS claim rule to leverage Regex and a common Active Directory naming convention for the security group to allow federation of all AWS accounts. Leverage third-party single sign-on as needed, and add it to the AD FS server.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://aws.amazon.com/ru/blogs/security/aws-federated-authentication-with-active-directory-federation-services-ad-fs/

QUESTION 153
In Amazon SNS, to send push notifications to mobile devices using Amazon SNS and ADM, you need to obtain the following, except:

A. Device token
B. Client ID
C. Registration ID
D. Client secret

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
To send push notifications to mobile devices using Amazon SNS and ADM, you need to obtain the following: Registration ID and Client secret.
Reference: http://docs.aws.amazon.com/sns/latest/dg/SNSMobilePushPrereq.html

QUESTION 154
A company is running multiple applications on Amazon EC2. Each application is deployed and managed by multiple business units. All applications are deployed on a single AWS account but on different virtual private clouds (VPCs). The company uses a separate VPC in the same account for test and development purposes.
Production applications suffered multiple outages when users accidentally terminated and modified resources that belonged to another business unit. A Solutions Architect has been asked to improve the availability of the company applications while allowing the Developers access to the resources they need.
Which option meets the requirements with the LEAST disruption?

A. Create an AWS account for each business unit. Move each business unit's instances to its own account and set up a federation to allow users to access their business unit's account.
B. Set up a federation to allow users to use their corporate credentials, and lock the users down to their own VPC. Use a network ACL to block each VPC from accessing other VPCs.
C. Implement a tagging policy based on business units. Create an IAM policy so that each user can terminate instances belonging to their own business units only.
D. Set up role-based access for each user and provide limited permissions based on individual roles and the services for which each user is responsible.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/IAM/latest/UserGuide/access_iam-tags.html

QUESTION 155
A user has set the IAM policy where it denies all requests if a request is not from IP 10.10.10.1/32. The other policy says allow all requests between 5 PM to 7 PM.
What will happen when a user is requesting access from IP 55.109.10.12/32 at 6 PM?

A. It will deny access
B. It is not possible to set a policy based on the time or IP
C. IAM will throw an error for policy conflict
D. It will allow access

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
When a request is made, the AWS IAM policy decides whether a given request should be allowed or denied. The evaluation logic follows these rules: By default, all requests are denied. (In general, requests made using the account credentials for resources in the account are always allowed.) An explicit allow policy overrides this default. An explicit deny policy overrides any allows.
In this case since there are explicit deny and explicit allow statements. Thus, the request will be denied since deny overrides allow.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_EvaluationLogic.html

QUESTION 156
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 two.)

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 Kinesis 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 157
When using string conditions within IAM, short versions of the available comparators can be used instead of the more verbose ones. streqi is the short version of the _____ string condition.

A. StringEqualsIgnoreCase
B. StringNotEqualsIgnoreCase
C. StringLikeStringEquals
D. StringNotEquals

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
When using string conditions within IAM, short versions of the available comparators can be used instead of the more verbose versions. For instance, streqi is the short version of StringEqualsIgnoreCase that checks for the exact match between two strings ignoring their case.
Reference: http://awsdocs.s3.amazonaws.com/SNS/20100331/sns-gsg-2010-03-31.pdf

QUESTION 158
A company needs to move its on-premises resources to AWS. The current environment consists of 100 virtual machines (VMs) with a total of 40 TB of storage. Most of the VMs can be taken offline because they support functions during business hours only, however, some are mission critical, so downtime must be minimized.
The administrator of the on-premises network provisioned 10 Mbps of internet bandwidth for the migration. The on-premises network throughput has reached capacity and would be costly to increase. A solutions architect must design a migration solution that can be performed within the next 3 months.
Which method would fulfill these requirements?

A. Set up a 1 Gbps AWS Direct Connect connection. Then, provision a private virtual interface, and use AWS Server Migration Service (SMS) to migrate the VMs into Amazon EC2.
B. Use AWS Application Discovery Service to assess each application, and determine how to refactor and optimize each using AWS services or AWS Marketplace solutions.
C. Export the VMs locally, beginning with the most mission-critical servers first. Use AWS Transfer for SFTP to securely upload each VM to Amazon S3 after they are exported. Use VM Import/Export to import the VMs into Amazon EC2.
D. Migrate mission-critical VMs with AWS SMS. Export the other VMs locally and transfer them to Amazon S3 using AWS Snowball. Use VM Import/Export to import the VMs into Amazon EC2.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 159
You want to establish redundant VPN connections and customer gateways on your network by setting up a second VPN connection.
Which of the following will ensure that this functions correctly?

A. The customer gateway IP address for the second VPN connection must be publicly accessible.
B. The virtual gateway IP address for the second VPN connection must be publicly accessible.
C. The customer gateway IP address for the second VPN connection must use dynamic routes.
D. The customer gateway IP address for the second VPN connection must be privately accessible and be the same public IP address that you are using for the first VPN connection.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
To establish redundant VPN connections and customer gateways on your network, you would need to set up a second VPN connection. However, you must ensure that the customer gateway IP address for the second VPN connection is publicly accessible.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_VPN.html

QUESTION 160
A solutions architect has an operational workload deployed on Amazon EC2 instances in an Auto Scaling group. The VPC architecture spans two Availability Zones (AZ) with a subnet in each that the Auto Scaling group is targeting. The VPC is connected to an on-premises environment and connectivity cannot be interrupted.
The maximum size of the Auto Scaling group is 20 instances in service. The VPC IPv4 addressing is as follows:
VPC CIDR: 10.0.0.0/23
AZ1 subnet CIDR: 10.0.0.0/24
AZ2 subnet CIDR: 10.0.1.0/24
Since deployment, a third AZ has become available in the Region. The solutions architect wants to adopt the new AZ without adding additional IPv4 address space and without service downtime.
Which solution will meet these requirements?

A. Update the Auto Scaling group to use the AZ2 subnet only. Delete and re-create the AZ1 subnet using half the previous address space. Adjust the Auto Scaling group to also use the new AZ1 subnet. When the instances are healthy, adjust the Auto Scaling group to use the AZ1 subnet only. Remove the current AZ2 subnet. Create a new AZ2 subnet using the second half of the address space from the original AZ1 subnet. Create a new AZ3 subnet using half the original AZ2 subnet address space, then update the Auto Scaling group to target all three new subnets.
B. Terminate the EC2 instances in the AZ1 subnet. Delete and re-create the AZ1 subnet using half the address space. Update the Auto Scaling group to use this new subnet. Repeat this for the second AZ. Define a new subnet in AZ3, then update the Auto Scaling group to target all three new subnets.

C. Create a new VPC with the same IPv4 address space and define three subnets, with one for each AZ. Update the existing Auto Scaling group to target the new subnets in the new VPC.

D. Update the Auto Scaling group to use the AZ2 subnet only. Update the AZ1 subnet to have the previous address space. Adjust the Auto Scaling group to also use the AZ1 subnet again. When the instances are healthy, adjust the Auto Scaling group to use the AZ1 subnet only. Update the current AZ2 subnet and assign the second half of the address space from the original AZ1 subnet. Create a new AZ3 subnet using half the original AZ2 subnet address space, then update the Auto Scaling group to target all three new subnets.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 161
A company has created an OU in AWS Organizations for each of its engineering teams. Each OU owns multiple AWS accounts. The organization has hundreds of AWS accounts.
A solutions architect must design a solution so that each OU can view a breakdown of usage costs across its AWS accounts.
Which solution meets these requirements?

A. Create an AWS Cost and Usage Report (CUR) for each OU by using AWS Resource Access Manager. Allow each team to visualize the CUR through an Amazon QuickSight dashboard.

B. Create an AWS Cost and Usage Report (CUR) from the AWS Organizations management account. Allow each team to visualize the CUR through an Amazon QuickSight dashboard.

C. Create an AWS Cost and Usage Report (CUR) in each AWS Organizations member account. Allow each team to visualize the CUR through an Amazon QuickSight dashboard.

D. Create an AWS Cost and Usage Report (CUR) by using AWS Systems Manager. Allow each team to visualize the CUR through Systems Manager OpsCenter dashboards.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/premiumsupport/knowledge-center/quicksight-cost-usage-report/

QUESTION 162
An organization has setup RDS with VPC. The organization wants RDS to be accessible from the internet. Which of the below mentioned configurations is not required in this scenario?

A. The organization must enable the parameter in the console which makes the RDS instance publicly accessible.
B. The organization must allow access from the internet in the RDS VPC security group,
C. The organization must setup RDS with the subnet group which has an external IP.
D. The organization must enable the VPC attributes DNS hostnames and DNS resolution.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources, such as RDS into a virtual network that the user has defined. Subnets are segments of a VPC's IP address range that the user can designate to a group of VPC resources based on security and operational needs. A DB subnet group is a collection of subnets (generally private) that the user can create in a VPC and which the user assigns to the RDS DB instances. A DB subnet group allows the user to specify a particular VPC when creating DB instances. If the RDS instance is required to be accessible from the internet:
The organization must setup that the RDS instance is enabled with the VPC attributes, DNS hostnames and DNS resolution.
The organization must enable the parameter in the console which makes the RDS instance publicly accessible. The organization must allow access from the internet in the RDS VPC security group.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.html

QUESTION 163
A user is creating a snapshot of an EBS volume. Which of the below statements is incorrect in relation to the creation of an EBS snapshot?

A. Its incremental
B. It is a point in time backup of the EBS volume
C. It can be used to create an AMI
D. It is stored in the same AZ as the volume

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The EBS snapshots are a point in time backup of the EBS volume. It is an incremental snapshot, but is always specific to the region and never specific to a single AZ. Hence the statement "It is stored in the same AZ as the volume" is incorrect.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSSnapshots.html

QUESTION 164
You want to use Amazon Redshift and you are planning to deploy dw1.8xlarge nodes. What is the minimum amount of nodes that you need to deploy with this kind of configuration?

A. 1
B. 4
C. 3
D. 2

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
For a single-node configuration in Amazon Redshift, the only option available is the smallest of the two options. The 8XL extra-large nodes are only available in a multi-node configuration.
Reference: http://docs.aws.amazon.com/redshift/latest/mgmt/working-with-clusters.html

QUESTION 165
A company is operating a large customer service call center, and stores and processes call recordings with a custom application. Approximately 2% of the call recordings are transcribed by an offshore team for quality assurance purposes.
These recordings take up to 72 hours to be transcribed. The recordings are stored on an NFS share before they are archived to an offsite location after 90 days.
The company uses Linux servers for processing the call recordings and managing the transcription queue. There is also a web application for the quality assurance staff to review and score call recordings.
The company plans to migrate the system to AWS to reduce storage costs and the time required to transcribe calls.
Which set of actions should be taken to meet the company's objectives?

A. Upload the call recordings to Amazon S3 from the call center. Set up an S3 lifecycle policy to move the call recordings to Amazon S3 Glacier after 90 days. Use an AWS Lambda trigger to transcribe the call recordings with Amazon Transcribe.
Use Amazon S3, Amazon API Gateway, and Lambda to host the review and scoring application.
B. Upload the call recordings to Amazon S3 from the call center. Set up an S3 lifecycle policy to move the call recordings to Amazon S3 Glacier after 90 days. Use an AWS Lambda trigger to transcribe the call recordings with Amazon Mechanical Turk. Use Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer to host the review and scoring application.
C. Use Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer to host the review and scoring application. Upload the call recordings to this application from the call center and store them on an Amazon EFS mount point. Use AWS Backup to archive the call recordings after 90 days. Transcribe the call recordings with Amazon Transcribe.
D. Upload the call recordings to Amazon S3 from the call center and put the object key in an Amazon SQS queue. Set up an S3 lifecycle policy to move the call recordings to Amazon S3 Glacier after 90 days. Use Amazon EC2 instances in an Auto Scaling group to send the recordings to Amazon Mechanical Turk for

transcription. Use the number of objects in the queue as the scaling metric. Use Amazon S3, Amazon API Gateway, and AWS Lambda to host the review and scoring application.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 166
A company is developing a messaging application that is based on a microservices architecture. A separate team develops each microservice by using Amazon Elastic Container Service (Amazon ECS). The teams deploy the microservices multiple times daily by using AWS CloudFormation and AWS CodePipeline. The application recently grew in size and complexity. Each service operates correctly on its own during development, but each service produces error messages when it has to interact with other services in production. A solutions architect must improve the application's availability.
Which solution will meet these requirements with the LEAST amount of operational overhead?

A. Add an extra stage to CodePipeline for each service. Use the extra stage to deploy each service to a test environment. Test each service after deployment to make sure that no error messages occur.
B. Add an AWS::CodeDeployBlueGreen Transform section and Hook section to the template to enable blue/green deployments by using AWS CodeDeploy in CloudFormation. Configure the template to perform ECS blue/green deployments in production.
C. Add an extra stage to CodePipeline for each service. Use the extra stage to deploy each service to a test environment. Write integration tests for each service. Run the tests automatically after deployment.
D. Use an ECS DeploymentConfiguration parameter in the template to configure AWS CodeDeploy to perform a rolling update of the service. Use a CircuitBreaker property to roll back the deployment if any error occurs during deployment.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/blogs/devops/using-aws-codepipeline-for-deploying-container-images-to-microservicesarchitecture-involving-aws-lambda-functions/

QUESTION 167
What types of identities do Amazon Cognito identity pools support?

A. They support both authenticated and unauthenticated identities.
B. They support only unauthenticated identities.

C. They support neither authenticated nor unauthenticated identities.
D. They support only authenticated identities.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Amazon Cognito identity pools support both authenticated and unauthenticated identities. Authenticated identities belong to users who are authenticated by a public login provider or your own backend authentication process. Unauthenticated identities typically belong to guest users.
Reference: http://docs.aws.amazon.com/cognito/devguide/identity/identity-pools/

QUESTION 168
A user has created a VPC with public and private subnets using the VPC wizard. The VPC has CIDR 20.0.0.0/16. The private subnet uses CIDR 20.0.0.0/24. The NAT instance ID is i-a12345.
Which of the below mentioned entries are required in the main route table attached with the private subnet to allow instances to connect with the internet?

A. Destination: 20.0.0.0/0 and Target: 80
B. Destination: 20.0.0.0/0 and Target: i-a12345
C. Destination: 20.0.0.0/24 and Target: i-a12345
D. Destination: 0.0.0.0/0 and Target: i-a12345

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A user can create a subnet with VPC and launch instances inside that subnet. If the user has created a public private subnet, the instances in the public subnet can receive inbound traffic directly from the Internet, whereas the instances in the private subnet cannot. If these subnets are created with Wizard, AWS will create two route tables and attach to the subnets. The main route table will have the entry "Destination: 0.0.0.0/0 and Target: i-a12345", which allows all the instances in the private subnet to connect to the internet using NAT.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario2.html

QUESTION 169
A company is migrating its applications to AWS. The applications will be deployed to AWS accounts owned by business units. The company has several teams of developers who are responsible for the development and maintenance of all applications. The company is expecting rapid growth in the number of users.
The company's chief technology officer has the following requirements:
Developers must launch the AWS infrastructure using AWS CloudFormation.

Developers must not be able to create resources outside of CloudFormation. The solution must be able to scale to hundreds of AWS accounts.
Which of the following would meet these requirements? (Choose two.)

A. Using CloudFormation, create an IAM role that can be assumed by CloudFormation that has permissions to create all the resources the company needs. Use CloudFormation StackSets to deploy this template to each AWS account.
B. In a central account, create an IAM role that can be assumed by developers, and attach a policy that allows interaction with CloudFormation. Modify the AssumeRolePolicyDocument action to allow the IAM role to be passed to CloudFormation.
C. Using CloudFormation, create an IAM role that can be assumed by developers, and attach policies that allow interaction with and passing a role to CloudFormation. Attach an inline policy to deny access to all other AWS services. Use CloudFormation StackSets to deploy this template to each AWS account.
D. Using CloudFormation, create an IAM role for each developer, and attach policies that allow interaction with CloudFormation. Use CloudFormation StackSets to deploy this template to each AWS account.
E. In a central AWS account, create an IAM role that can be assumed by CloudFormation that has permissions to create the resources the company requires. Create a CloudFormation stack policy that allows the IAM role to manage resources. Use CloudFormation StackSets to deploy the CloudFormation stack policy to each AWS account.

Correct Answer: CE
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_boundaries.html

QUESTION 170
A manufacturing company is growing exponentially and has secured funding to improve its IT infrastructure and ecommerce presence. The company's ecommerce platform consists of:
Static assets primarily comprised of product images stored in Amazon S3.
Amazon DynamoDB tables that store product information, user information, and order information. Web servers containing the application's front-end behind Elastic Load Balancers.
The company wants to set up a disaster recovery site in a separate Region.
Which combination of actions should the solutions architect take to implement the new design while meeting all the requirements? (Choose three.)

A. Enable Amazon Route 53 health checks to determine if the primary site is down, and route traffic to the disaster recovery site if there is an issue.
B. Enable Amazon S3 cross-Region replication on the buckets that contain static assets.
C. Enable multi-Region targets on the Elastic Load Balancer and target Amazon EC2 instances in both Regions.
D. Enable DynamoDB global tables to achieve a multi-Region table replication.
E. Enable Amazon CloudWatch and create CloudWatch alarms that route traffic to the disaster recovery site when application latency exceeds the desired threshold.

F. Enable Amazon S3 versioning on the source and destination buckets containing static assets to ensure there is a rollback version available in the event of data corruption.

Correct Answer: AEF
Section: (none)
Explanation

Explanation/Reference:

QUESTION 171
DynamoDB uses only as a transport protocol, not as a storage format.

A. WDDX
B. XML
C. SGML
D. JSON

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
DynamoDB uses JSON only as a transport protocol, not as a storage format. The AWS SDKs use JSON to send data to DynamoDB, and DynamoDB responds with JSON, but DynamoDB does not store data persistently in JSON format.
Reference: http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Programming.LowLev elAPI.html

QUESTION 172
A company is using AWS CloudFormation to deploy its infrastructure. The company is concerned that, if a production CloudFormation stack is deleted, important data stored in Amazon RDS databases or Amazon EBS volumes might also be deleted.
How can the company prevent users from accidentally deleting data in this way?

A. Modify the CloudFormation templates to add a DeletionPolicy attribute to RDS and EBS resources.
B. Configure a stack policy that disallows the deletion of RDS and EBS resources.
C. Modify IAM policies to deny deleting RDS and EBS resources that are tagged with an “aws:cloudformation:stack-name” tag.
D. Use AWS Config rules to prevent deleting RDS and EBS resources.

Correct Answer: A

Section: (none)
Explanation

Explanation/Reference:
Explanation:
With the DeletionPolicy attribute you can preserve or (in some cases) backup a resource when its stack is deleted. You specify a DeletionPolicy attribute for each resource that you want to control. If a resource has no DeletionPolicy attribute, AWS CloudFormation deletes the resource by default. To keep a resource when its stack is deleted, specify Retain for that resource. You can use retain for any resource. For example, you can retain a nested stack, Amazon S3 bucket, or EC2 instance so that you can continue to use or modify those resources after you delete their stacks.
Reference:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-deletionpolicy.html

QUESTION 173
A company is running a containerized application in the AWS Cloud. The application is running by using Amazon Elastic Container Service (Amazon ECS) on a set of Amazon EC2 instances. The EC2 instances run in an Auto Scaling group.
The company uses Amazon Elastic Container Registry (Amazon ECR) to store its container images. When a new image version is uploaded, the new image version receives a unique tag.
The company needs a solution that inspects new image versions for common vulnerabilities and exposures. The solution must automatically delete new image tags that have Critical or High severity findings. The solution also must notify the development team when such a deletion occurs.
Which solution meets these requirements?

A. Configure scan on push on the repository. Use Amazon EventBridge (Amazon CloudWatch Events) to invoke an AWS Step Functions state machine when a scan is complete for images that have Critical or High severity findings. Use the Step Functions state machine to delete the image tag for those images and to notify the development team through Amazon Simple Notification Service (Amazon SNS).
B. Configure scan on push on the repository. Configure scan results to be pushed to an Amazon Simple Queue Service (Amazon SQS) queue. Invoke an AWS Lambda function when a new message is added to the SQS queue. Use the Lambda function to delete the image tag for images that have Critical or High severity findings. Notify the development team by using Amazon Simple Email Service (Amazon SES).
C. Schedule an AWS Lambda function to start a manual image scan every hour. Configure Amazon EventBridge (Amazon CloudWatch Events) to invoke another Lambda function when a scan is complete. Use the second Lambda function to delete the image tag for images that have Critical or High severity findings. Notify the development team by using Amazon Simple Notification Service (Amazon SNS).
D. Configure periodic image scan on the repository. Configure scan results to be added to an Amazon Simple Queue Service (Amazon SQS) queue. Invoke an AWS Step Functions state machine when a new message is added to the SQS queue. Use the Step Functions state machine to delete the image tag for images that have Critical or High severity findings.
Notify the development team by using Amazon Simple Email Service (Amazon SES).

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

Explanation:
Reference: https://docs.aws.amazon.com/AmazonCloudWatch/latest/events/cwe-ug.pdf

QUESTION 174
You need to develop and run some new applications on AWS and you know that Elastic Beanstalk and CloudFormation can both help as a deployment mechanism for a broad range of AWS resources.
Which of the following is TRUE statements when describing the differences between Elastic Beanstalk and CloudFormation?

A. AWS Elastic Beanstalk introduces two concepts: The template, a JSON or YAML-format, text- based file
B. Elastic Beanstalk supports AWS CloudFormation application environments as one of the AWS resource types.
C. Elastic Beanstalk automates and simplifies the task of repeatedly and predictably creating groups of related resources that power your applications. CloudFormation does not.
D. You can design and script custom resources in CloudFormation

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
These services are designed to complement each other. AWS Elastic Beanstalk provides an environment to easily deploy and run applications in the cloud. It is integrated with developer tools and provides a one-stop experience for you to manage the lifecycle of your applications. AWS CloudFormation is a convenient provisioning mechanism for a broad range of AWS resources. It supports the infrastructure needs of many different types of applications such as existing enterprise applications, legacy applications, applications built using a variety of AWS resources and container-based solutions (including those built using AWS Elastic Beanstalk). AWS CloudFormation supports Elastic Beanstalk application environments as one of the AWS resource types. This allows you, for example, to create and manage an AWS Elastic Beanstalk- hosted application along with an RDS database to store the application data. In addition to RDS instances, any other supported AWS resource can be added to the group as well.
Reference: https://aws.amazon.com/cloudformation/faqs

QUESTION 175
You create an Amazon Elastic File System (EFS) file system and mount targets for the file system in your Virtual Private Cloud (VPC). Identify the initial permissions you can grant to the group root of your file system.

A. write-execute-modify
B. read-execute
C. read-write-modify
D. read-write

Correct Answer: B

Section: (none)
Explanation

Explanation/Reference:
Explanation:
In Amazon EFS, when a file system and mount targets are created in your VPC, you can mount the remote file system locally on your Amazon Elastic Compute Cloud (EC2) instance. You can grant permissions to the users of your file system.
The initial permissions mode allowed for Amazon EFS are: read-write-execute permissions to the owner root read-execute permissions to the group root read-execute permissions to others
Reference: http://docs.aws.amazon.com/efs/latest/ug/accessing-fs-nfs-permissions.html

QUESTION 176
A company hosts a game player-matching service on a public facing, physical, on-premises instance that all users are able to access over the internet. All traffic to the instance uses UDP. The company wants to migrate the service to AWS and provide a high level of security. A solutions architect needs to design a solution for the player-matching service using AWS.
Which combination of steps should the solutions architect take to meet these requirements? (Choose three.)

A. Use a Network Load Balancer (NLB) in front of the player-matching instance. Use a friendly DNS entry in Amazon Route 53 pointing to the NLB's Elastic IP address.
B. Use an Application Load Balancer (ALB) in front of the player-matching instance. Use a friendly DNS entry in Amazon Route 53 pointing to the ALB's internet-facing fully qualified domain name (FQDN).
C. Define an AWS WAF rule to explicitly drop non-UDP traffic, and associate the rule with the load balancer.
D. Configure a network ACL rule to block all non-UDP traffic. Associate the network ACL with the subnets that hold the load balancer instances.
E. Use Amazon CloudFront with an Elastic Load Balancer as an origin.
F. Enable AWS Shield Advanced on all public-facing resources.

Correct Answer: BDF
Section: (none)
Explanation

Explanation/Reference:

QUESTION 177
An organization, which has the AWS account ID as 999988887777, has created 50 IAM users. All the users are added to the same group ABC. If the organization has enabled that each IAM user can login with the AWS console, which AWS login URL will the IAM users use??

A. https://999988887777.aws.amazon.com/ABC/
B. https://signin.aws.amazon.com/ABC/

C. https://ABC.signin.aws.amazon.com/999988887777/console/
D. https://999988887777.signin.aws.amazon.com/console/

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. Once the organization has created the IAM users, they will have a separate AWS console URL to login to the AWS console. The console login URL for the IAM user will be https:// AWS_Account_ID.signin.aws.amazon.com/console/. It uses only the AWS account ID and does not depend on the group or user ID.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/AccountAlias.html

QUESTION 178
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 CIoudFront 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:
Explanation:
You can have CloudFront sit in front of your on-prem web environment, via a custom origin (the origin doesn't have to be in AWS). This would protect against unexpected bursts in traffic by letting CloudFront handle the traffic that it can out of cache, thus hopefully removing some of the load from your on-prem web

servers.

QUESTION 179
A company's data center is connected to the AWS Cloud over a minimally used 10 Gbps AWS Direct Connect connection with a private virtual interface to its virtual private cloud (VPC). The company internet connection is 200 Mbps, and the company has a 150 TB dataset that is created each Friday. The data must be transferred and available in Amazon S3 on Monday morning.
Which is the LEAST expensive way to meet the requirements while allowing for data transfer growth?

A. Order two 80 TB AWS Snowball appliances. Offload the data to the appliances and ship them to AWS. AWS will copy the data from the Snowball appliances to Amazon S3.
B. Create a VPC endpoint for Amazon S3. Copy the data to Amazon S3 by using the VPC endpoint, forcing the transfer to use the Direct Connect connection.
C. Create a VPC endpoint for Amazon S3. Set up a reverse proxy farm behind a Classic Load Balancer in the VPCopy the data to Amazon S3 using the proxy.
D. Create a public virtual interface on a Direct Connect connection, and copy the data to Amazon S3 over the connection.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 180
You want to define permissions for a role in an IAM policy. Which of the following configuration formats should you use?

A. An XML document written in the IAM Policy Language
B. An XML document written in a language of your choice
C. A JSON document written in the IAM Policy Language
D. JSON document written in a language of your choice

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
You define the permissions for a role in an IAM policy. An IAM policy is a JSON document written in the IAM Policy Language.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_terms-and-concepts.html

QUESTION 181

In Amazon ElastiCache, the default cache port is:

A. for Memcached 11210 and for Redis 6380.
B. for Memcached 11211 and for Redis 6380.
C. for Memcached 11210 and for Redis 6379.
D. for Memcached 11211 and for Redis 6379.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In Amazon ElastiCache, you can specify a new port number for your cache cluster, which by default is 11211 for Memcached and 6379 for Redis.
Reference: http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/GettingStarted.AuthorizeAccess.html

QUESTION 182
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 onpremises 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 two.)

A. Develop an identity broker that authenticates against IAM security Token service to assume a IAM role in order to get temporary AWS security credentials The application calls the identity broker to get AWS temporary security credentials with access to the appropriate S3 bucket.
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:
Explanation:

Imagine that in your organization, you want to provide a way for users to copy data from their computers to a backup folder.
You build an application that users can run on their computers. On the back end, the application reads and writes objects in an S3 bucket. Users don't have direct access to AWS. Instead, the application communicates with an identity provider (IdP) to authenticate the user. The IdP gets the user information from your organization's identity store (such as an LDAP directory) and then generates a SAML assertion that includes authentication and authorization information about that user.
The application then uses that assertion to make a call to the AssumeRoleWithSAML API to get temporary security credentials. The app can then use those credentials to access a folder in the S3 bucket that's specific to the user.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_saml.html

QUESTION 183
A company that provisions job boards for a seasonal workforce is seeing an increase in traffic and usage. The backend services run on a pair of Amazon EC2 instances behind an Application Load Balancer with Amazon DynamoDB as the datastore. Application read and write traffic is slow during peak seasons.
Which option provides a scalable application architecture to handle peak seasons with the LEAST development effort?

A. Migrate the backend services to AWS Lambda. Increase the read and write capacity of DynamoDB
B. Migrate the backend services to AWS Lambda. Configure DynamoDB to use global tables
C. Use Auto Scaling groups for the backend services. Use DynamoDB auto scaling
D. Use Auto Scaling groups for the backend services. Use Amazon Simple Queue Service (Amazon SQS) and an AWS Lambda function to write to DynamoDB

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 184
Which of the following is not included in the metrics sent from Billing to Amazon CloudWatch?

A. Recurring fees for AWS products and services
B. Total AWS charges
C. One-time charges and refunds
D. Usage charges for AWS products and services

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

Explanation:
Usage charges and recurring fees for AWS products and services are included in the metrics sent from Billing to Amazon CloudWatch. You will have a metric for total AWS charges, as well as one additional metric for each AWS product or service that you use. However, one-time charges and refunds are not included.
Reference:
https://aws.amazon.com/blogs/aws/monitor-estimated-costs-using-amazon-cloudwatch-billing-metrics-and-alarms

QUESTION 185
In AWS IAM, which of the following predefined policy condition keys checks how long ago (in seconds) the MFA-validated security credentials making the request were issued using multi- factor authentication (MFA)?

A. aws:MultiFactorAuthAge
B. aws:MultiFactorAuthLast
C. aws:MFAAge
D. aws:MultiFactorAuthPrevious

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation: aws:MultiFactorAuthAge is one of the predefined keys provided by AWS that can be included within a Condition element of an IAM policy. The key allows to check how long ago (in seconds) the MFAvalidated security credentials making the request were issued using Multi-Factor Authentication (MFA).
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_ElementDescriptions.html

QUESTION 186
A company is planning on hosting its ecommerce platform on AWS using a multi-tier web application designed for a NoSQL database. The company plans to use the us-west-2 Region as its primary Region. The company wants to ensure that copies of the application and data are available in second Region, us-west-1, for disaster recovery. The company wants to keep the time to fail over as low as possible. Failing back to the primary Region should be possible without administrative interaction after the primary service is restored.
Which design should the solutions architect use?

A. Use AWS CloudFormation StackSets to create the stacks in both Regions with Auto Scaling groups for the web and application tiers. Asynchronously replicate static content between Regions using Amazon S3 cross-Region replication. Use an Amazon Route 53 DNS failover routing policy to direct users to the secondary site in us-west-1 in the event of an outage.
Use Amazon DynamoDB global tables for the database tier.
B. Use AWS CloudFormation StackSets to create the stacks in both Regions with Auto Scaling groups for the web and application tiers. Asynchronously replicate static content between Regions using Amazon S3 cross-Region replication. Use an Amazon Route 53 DNS failover routing policy to direct users to the secondary site in us-west-1 in the event of an outage Deploy an Amazon Aurora global database for the database tier.
C. Use AWS Service Catalog to deploy the web and application servers in both Regions Asynchronously replicate static content between the two Regions using Amazon S3 cross-Region replication. Use Amazon Route 53 health checks to identify a primary Region failure and update the public DNS entry listing to the

secondary Region in the event of an outage.
Use Amazon RDS for MySQL with crossRegion replication for the database tier.

D. Use AWS CloudFormation StackSets to create the stacks in both Regions using Auto Scaling groups for the web and application tiers. Asynchronously replicate static content between Regions using Amazon S3 cross-Region replication. Use Amazon CloudFront with static files in Amazon S3, and multi-Region origins for the front-end web tier. Use Amazon DynamoDB tables in each Region with scheduled backups to Amazon S3.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 187
In IAM, which of the following is true of temporary security credentials?

A. Once you issue temporary security credentials, they cannot be revoked.
B. None of these are correct.
C. Once you issue temporary security credentials, they can be revoked only when the virtual MFA device is used.
D. Once you issue temporary security credentials, they can be revoked.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Temporary credentials in IAM are valid throughout their defined duration of time and hence can't be revoked. However, because permissions are evaluated each time an AWS request is made using the credentials, you can achieve the effect of revoking the credentials by changing the permissions for the credentials even after they have been issued.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_temp_control-access_disable-perms.html

QUESTION 188
A company with global offices has a single 1 Gbps AWS Direct Connect connection to a single AWS Region. The company's on-premises network uses the connection to communicate with the company's resources in the AWS Cloud. The connection has a single private virtual interface that connects to a single VPC. A solutions architect must implement a solution that adds a redundant Direct Connect connection in the same Region. The solution also must provide connectivity to other Regions through the same pair of Direct Connect connections as the company expands into other Regions.
Which solution meets these requirements?

A. Provision a Direct Connect gateway. Delete the existing private virtual interface from the existing connection. Create the second Direct Connect connection.

Create a new private virtual interface on each connection, and connect both private virtual interfaces to the Direct Connect gateway. Connect the Direct Connect gateway to the single VPC.

B. Keep the existing private virtual interface. Create the second Direct Connect connection. Create a new private virtual interface on the new connection, and connect the new private virtual interface to the single VPC.

C. Keep the existing private virtual interface. Create the second Direct Connect connection. Create a new public virtual interface on the new connection and connect the new public virtual interface to the single VPC.

D. Provision a transit gateway. Delete the existing private virtual interface from the existing connection. Create the second Direct Connect connection. Create a new private virtual interface on each connection, and connect both private virtual interfaces to the transit gateway. Associate the transit gateway with the single VPC.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/directconnect/latest/UserGuide/WorkingWithVirtualInterfaces.html

QUESTION 189
A user has created a VPC with two subnets: one public and one private. The user is planning to run the patch update for the instances in the private subnet. How can the instances in the private subnet connect to the internet?

A. The private subnet can never connect to the internet

B. Use NAT with an elastic IP

C. Use the internet gateway with a private IP

D. Allow outbound traffic in the security group for port 80 to allow internet updates

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. If the user has created two subnets (one private and one public), they would need a Network Address Translation (NAT) instance with the elastic IP address. This enables the instances in the private subnet to send requests to the internet (for example, to perform software updates).
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario2.html

QUESTION 190

A company has an internal application running on AWS that is used to track and process shipments in the company's warehouse. Currently, after the system receives an order, it emails the staff the information needed to ship a package. Once the package is shipped, the staff replies to the email and the order is marked as shipped.
The company wants to stop using email in the application and move to a serverless application model.
Which architecture solution meets these requirements?

A. Use AWS Batch to configure the different tasks required to ship a package. Have AWS Batch trigger an AWS Lambda function that creates and prints a shipping label. Once that label is scanned, as it leaves the warehouse, have another Lambda function move the process to the next step in the AWS Batch job.
B. When a new order is created, store the order information in Amazon SQS. Have AWS Lambda check the queue every 5 minutes and process any needed work. When an order needs to be shipped, have Lambda print the label in the warehouse.
Once the label has been scanned, as it leaves the warehouse, have an Amazon EC2 instance update Amazon SQS.
C. Update the application to store new order information in Amazon DynamoDB. When a new order is created, trigger an AWS Step Functions workflow, mark the orders as "in progress", and print a package label to the warehouse. Once the label has been scanned and fulfilled, the application will trigger an AWS Lambda function that will mark the order as shipped and complete the workflow.
D. Store new order information in Amazon EFS. Have instances pull the new information from the NFS and send that information to printers in the warehouse. Once the label has been scanned, as it leaves the warehouse, have Amazon API Gateway call the instances to remove the order information from Amazon EFS.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 191
_____pricing offers significant savings over the normal price of DynamoDB provisioned throughput capacity.

A. Discount Voucher
B. Reserved Capacity
C. Discount Service
D. Reserved Point

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

Explanation:
Reserved Capacity pricing offers significant savings over the normal price of DynamoDB provisioned throughput capacity.
When you buy Reserved Capacity, you pay a one-time upfront fee and commit to paying for a minimum usage level, at the hourly rates indicated above, for the duration of the Reserved Capacity term.
Reference: http://aws.amazon.com/dynamodb/pricing/

QUESTION 192
When you resize the Amazon RDS DB instance, Amazon RDS will perform the upgrade during the next maintenance window. If you want the upgrade to be performed now, rather than waiting for the maintenance window, specify the option.

A. ApplyNow
B. ApplySoon
C. ApplyThis
D. ApplyImmediately

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.DBInstance.Modifying.html

QUESTION 193
A company is using AWS Organizations to manage 15 AWS accounts. A solutions architect wants to run advanced analytics on the company's cloud expenditures. The cost data must be gathered and made available from an analytics account. The analytics application runs in a VPC and must receive the raw cost data each night to run the analytics.
The solutions architect has decided to use the Cost Explorer API to fetch the raw data and store the data in Amazon S3 in JSON format. Access to the raw cost data must be restricted to the analytics application. The solutions architect has already created an AWS Lambda function to collect data by using the Cost Explorer API.
Which additional actions should the solutions architect take to meet these requirements?

A. Create an IAM role in the Organizations master account with permissions to use the Cost Explorer API, and establish trust between the role and the analytics account. Update the Lambda function role and add sts:AssumeRole permissions. Assume the role in the master account from the Lambda function code by using the AWS Security Token Service (AWS STS) AssumeRole API call. Create a gateway endpoint for Amazon S3 in the analytics VPC. Create an S3 bucket policy that allows access only from the S3 endpoint.
B. Create an IAM role in the analytics account with permissions to use the Cost Explorer API. Update the Lambda function and assign the new role. Create a gateway endpoint for Amazon S3 in the analytics VPC. Create an S3 bucket policy that allows access only from the analytics VPC by using the aws:SourceVpc condition.
C. Create an IAM role in the Organizations master account with permissions to use the Cost Explorer API, and establish trust between the role and the analytics

account. Update the Lambda function role and add sts:AssumeRole permissions. Assume the role in the master account from the Lambda function code by using the AWS Security Token Service (AWS STS) AssumeRole API call. Create an interface endpoint for Amazon S3 in the analytics VPCreate an S3 bucket policy that allows access only from the analytics VPC private CIDR range by using the aws:SourceIp condition.

D. Create an IAM role in the analytics account with permissions to use the Cost Explorer API. Update the Lambda function and assign the new role. Create an interface endpoint for Amazon S3 in the analytics VPC. Create an S3 bucket policy that allows access only from the S3 endpoint.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 194
A company uses AWS Organizations with a single OU named Production to manage multiple accounts. All accounts are members of the Production OU. Administrators use deny list SCPs in the root of the organization to manage access to restricted services.
The company recently acquired a new business unit and invited the new unit's existing AWS account to the organization.
Once onboarded, the administrators of the new business unit discovered that they are not able to update existing AWS Config rules to meet the company's policies.
Which option will allow administrators to make changes and continue to enforce the current policies without introducing additional long-term maintenance?

A. Remove the organization's root SCPs that limit access to AWS Config. Create AWS Service Catalog products for the company's standard AWS Config rules and deploy them throughout the organization, including the new account.

B. Create a temporary OU named Onboarding for the new account. Apply an SCP to the Onboarding OU to allow AWS Config actions. Move the new account to the Production OU when adjustments to AWS Config are complete.

C. Convert the organization's root SCPs from deny list SCPs to allow list SCPs to allow the required services only. Temporally apply an SCP to the organization's root that allows AWS Config actions for principals only in the new account.

D. Create a temporary OU named Onboarding for the new account. Apply an SCP to the Onboarding OU to allow AWS Config actions. Move the organization's root SCP to the Production OU. Move the new account to the Production OU when adjustments to AWS Config are complete.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 195
What bandwidths do AWS Direct Connect currently support?

A. 10Mbps and 100Mbps
B. 10Gbps and 100Gbps
C. 100Mbps and 1Gbps
D. 1Gbps and 10 Gbps

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
AWS Direct Connection currently supports 1Gbps and 10 Gbps.
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.html

QUESTION 196
The company Security team requires that all data uploaded into an Amazon S3 bucket must be encrypted. The encryption keys must be highly available and the company must be able to control access on a per-user basis, with different users having access to different encryption keys.
Which of the following architectures will meet these requirements? (Choose two.)

A. Use Amazon S3 server-side encryption with Amazon S3-managed keys. Allow Amazon S3 to generate an AWS/S3 master key, and use IAM to control access to the data keys that are generated.
B. Use Amazon S3 server-side encryption with AWS KMS-managed keys, create multiple customer master keys, and use key policies to control access to them.
C. Use Amazon S3 server-side encryption with customer-managed keys, and use AWS CloudHSM to manage the keys. Use CloudHSM client software to control access to the keys that are generated.
D. Use Amazon S3 server-side encryption with customer-managed keys, and use two AWS CloudHSM instances configured in high-availability mode to manage the keys. Use the CloudHSM client software to control access to the keys that are generated.
E. Use Amazon S3 server-side encryption with customer-managed keys, and use two AWS CloudHSM instances configured in high-availability mode to manage the keys. Use IAM to control access to the keys that are generated in CloudHSM.

Correct Answer: BC
Section: (none)
Explanation

Explanation/Reference:

QUESTION 197
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

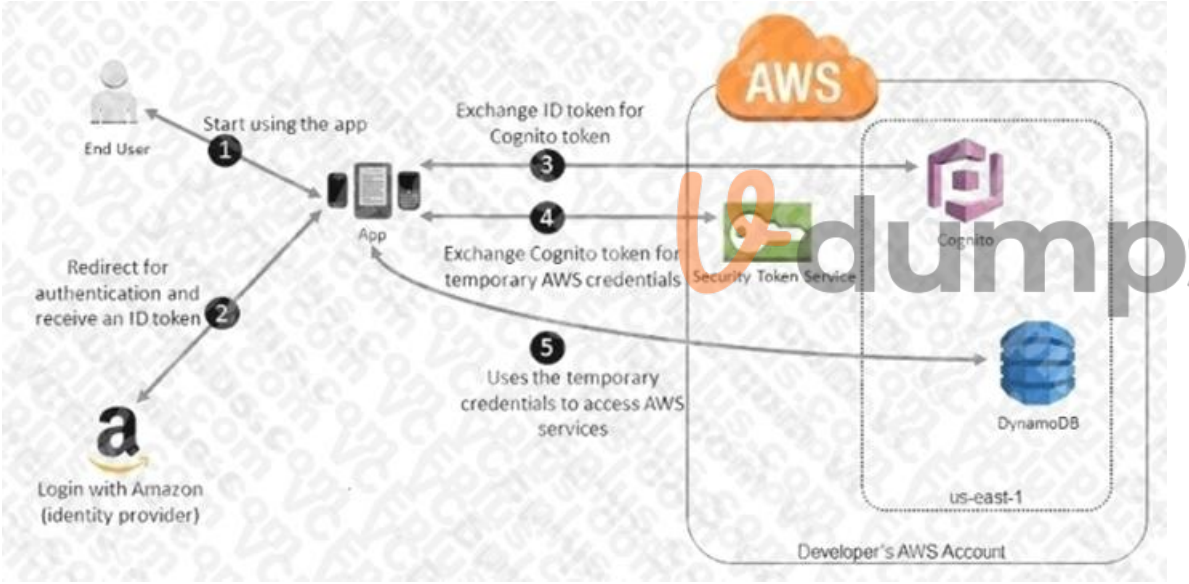

A customer starts your app on a mobile device. The app asks the user to sign in.
The app uses Login with Amazon resources to accept the user's credentials.
The app uses Cognito APIs to exchange the Login with Amazon ID token for a Cognito token.
The app requests temporary security credentials from AWS 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 198
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 rearview 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 199
Does Autoscaling automatically assign tags to resources?

A. No, not unless they are configured via API.

B. Yes, it does.
C. Yes, by default.
D. No, it does not.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Tags don't have any semantic meaning to Amazon EC2 and are interpreted strictly as a string of characters.
Tags are assigned automatically to the instances created by an Auto Scaling group. Auto Scaling adds a tag to the instance with a key of aws: autoscaling:groupName and a value of the name of the Auto Scaling group.
Reference: http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/Using_Tags.html

QUESTION 200
A user has launched an EBS optimized instance with EC2. Which of the below mentioned options is the correct statement?

A. It provides additional dedicated capacity for EBS IO
B. The attached EBS will have greater storage capacity
C. The user will have a PIOPS based EBS volume
D. It will be launched on dedicated hardware in VPC

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
An Amazon EBS-optimized instance uses an optimized configuration stack and provides additional, dedicated capacity for the Amazon EBS I/O. This optimization provides the best performance for the user's Amazon EBS volumes by minimizing contention between the Amazon EBS I/O and other traffic from the user's instance.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSOptimized.html

QUESTION 201
A company needs to run a software package that has a license that must be run on the same physical host for the duration of its use. The software package is only going to be used for 90 days. The company requires patching and restarting of all instances every 30 days.
How can these requirements be met using AWS?

A. Run a dedicated instance with auto-placement disabled.
B. Run the instance on a dedicated host with Host Affinity set to Host.
C. Run an On-Demand Instance with a Reserved Instance to ensure consistent placement.
D. Run the instance on a licensed host with termination set for 90 days.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/how-dedicated-hosts-work.html

QUESTION 202
A company is running a workload that consists of thousands of Amazon EC2 instances. The workload is running in a VPC that contains several public subnets and private subnets. The public subnets have a route for 0.0.0.0/0 to an existing internet gateway. The private subnets have a route for 0.0.0.0/0 to an existing NAT gateway.
A solutions architect needs to migrate the entire fleet of EC2 instances to use IPv6. The EC2 instances that are in private subnets must not be accessible from the public internet.
What should the solutions architect do to meet these requirements?

A. Update the existing VPC, and associate a custom IPv6 CIDR block with the VPC and all subnets. Update all the VPC route tables, and add a route for ::/0 to the internet gateway.
B. Update the existing VPC, and associate an Amazon-provided IPv6 CIDR block with the VPC and all subnets. Update the VPC route tables for all private subnets, and add a route for ::/0 to the NAT gateway.
C. Update the existing VPC, and associate an Amazon-provided IPv6 CIDR block with the VPC and all subnets. Create an egress-only internet gateway. Update the VPC route tables for all private subnets, and add a route for ::/0 to the egress-only internet gateway.
D. Update the existing VPC, and associate a custom IPv6 CIDR block with the VPC and all subnets. Create a new NAT gateway, and enable IPv6 support. Update the VPC route tables for all private subnets, and add a route for ::/0 to the IPv6- enabled NAT gateway.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 203
A user wants to create a public subnet in VPC and launch an EC2 instance within it. The user has not selected the option to assign a public IP address while

launching the instance.
Which of the below mentioned statements is true with respect to this scenario?

A. The instance will always have a public DNS attached to the instance by default
B. The user would need to create a default route to IGW in subnet's route table and then attach an elastic IP to the instance to connect from the internet
C. The user can directly attach an elastic IP to the instance
D. The instance will never launch if the public IP is not assigned

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. When the user is launching an instance he needs to select an option which attaches a public IP to the instance. If the user has not selected the option to attach the public IP, then it will only have a private IP when launched. The user cannot connect to the instance from the internet. If the user wants an elastic IP to connect to the instance from the internet, he should create an internet gateway and assign an elastic IP to instance.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/GettingStartedGuide/LaunchInstance.html

QUESTION 204
You have set up a huge amount of network infrastructure in AWS and you now need to think about monitoring all of this. You decide CloudWatch will best fit your needs but you are unsure of the pricing structure and the limitations of CloudWatch.
Which of the following statements is TRUE in relation to the limitations of CloudWatch?

A. You get 10 CloudWatch metrics, 10 alarms, 1,000,000 API requests, and 1,000 Amazon SNS email notifications per customer per month for free.
B. You get 100 CloudWatch metrics, 100 alarms, 10,000,000 API requests, and 10,000 Amazon SNS email notifications per customer per month for free.
C. You get 10 CloudWatch metrics, 10 alarms, 1,000 API requests, and 100 Amazon SNS email notifications per customer per month for free.
D. You get 100 CloudWatch metrics, 100 alarms, 1,000,000 API requests, and 1,000 Amazon SNS email notifications per customer per month for free.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Amazon CloudWatch monitors your Amazon Web Services (AWS) resources and the applications you run on AWS in realtime.
You can use CloudWatch to collect and track metrics, which are the variables you want to measure for your resources and applications. CloudWatch has the following limits:

You get 10 CloudWatch metrics, 10 alarms, 1,000,000 API requests, and 1,000 Amazon SNS email notifications per customer per month for free. You can assign up to 10 dimensions per metric.
You can create up to 5000 alarms per AWS account. Metric data is kept for 2 weeks.
The size of a PutMetricData request is limited to 8KB for HTTP GET requests and 40KB for HTTP POST requests.
You can include a maximum of 20 MetricDatum items in one PutMetricData request. A MetricDatum can contain a single value or a StatisticSet representing many values.
Reference: http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/cloudwatch_limits.html

QUESTION 205
A company must deploy multiple independent instances of an application. The front-end application is internet accessible.
However, corporate policy stipulates that the backends are to be isolated from each other and the internet, yet accessible from a centralized administration server. The application setup should be automated to minimize the opportunity for mistakes as new instances are deployed.
Which option meets the requirements and MINIMIZES costs?

A. Use an AWS CloudFormation template to create identical IAM roles for each region. Use AWS CloudFormation StackSets to deploy each application instance by using parameters to customize for each instance, and use security groups to isolate each instance while permitting access to the central server.
B. Create each instance of the application IAM roles and resources in separate accounts by using AWS CloudFormation StackSets. Include a VPN connection to the VPN gateway of the central administration server.
C. Duplicate the application IAM roles and resources in separate accounts by using a single AWS CloudFormation template. Include VPC peering to connect the VPC of each application instance to a central VPC.
D. Use the parameters of the AWS CloudFormation template to customize the deployment into separate accounts. Include a NAT gateway to allow communication back to the central administration server.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 206
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:
Explanation:
Reference:
https://aws.amazon.com/blogs/security/writing-iam-policies-grant-access-to-user-specific-folders- in-an-amazon-s3-bucket/

QUESTION 207
How many g2.2xlarge on-demand instances can a user run in one region without taking any limit increase approval from AWS?

A. 20
B. 2
C. 5
D. 10

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Generally, AWS EC2 allows running 20 on-demand instances and 100 spot instances at a time. This limit can be increased by requesting at https:// aws.amazon.com/contact-us/ec2-request. Excluding certain types of instances, the limit is lower than mentioned above. For g2.2xlarge, the user can run only 5 on-demand instance at a time.
Reference: http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#limits_ec2

QUESTION 208
A company has a complex web application that leverages Amazon CloudFront for global scalability and performance. Over time, users report that the web application is slowing down.
The company's operations team reports that the CloudFront cache hit ratio has been dropping steadily. The cache metrics report indicates that query strings on some URLs are inconsistently ordered and are specified sometimes in mixed-case letters and sometimes in lowercase letters.
Which set of actions should the solutions architect take to increase the cache hit ratio as quickly possible?

A. Deploy a Lambda@Edge function to sort parameters by name and force them to be lowercase. Select the CloudFront viewer request trigger to invoke the

function.

B. Update the CloudFront distribution to disable caching based on query string parameters.

C. Deploy a reverse proxy after the load balancer to post process the emitted URLs in the application to force the URL strings to be lowercase.

D. Update the CloudFront distribution to specify case-insensitive query string processing.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 209
A solutions architect must implement a multi-Region architecture for an Amazon RDS for PostgreSQL database that supports a web application. The database launches from an AWS CloudFormation template that includes AWS services and features that are present in both the primary and secondary Regions.
The database is configured for automated backups, and it has an RTO of 15 minutes and an RPO of 2 hours. The web application is configured to use an Amazon Route 53 record to route traffic to the database.
Which combination of steps will result in a highly available architecture that meets all the requirements? (Choose two.)

A. Create a cross-Region read replica of the database in the secondary Region. Configure an AWS Lambda function in the secondary Region to promote the read replica during failover event.

B. In the primary Region, create a health check on the database that will invoke an AWS Lambda function when a failure is detected. Program the Lambda function to recreate the database from the latest database snapshot in the secondary Region and update the Route 53 host records for the database.

C. Create an AWS Lambda function to copy the latest automated backup to the secondary Region every 2 hours.

D. Create a failover routing policy in Route 53 for the database DNS record. Set the primary and secondary endpoints to the endpoints in each Region.

E. Create a hot standby database in the secondary Region. Use an AWS Lambda function to restore the secondary database to the latest RDS automatic backup in the event that the primary database fails.

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover-types.html

QUESTION 210
A user is thinking to use EBS PIOPS volume.

Which of the below mentioned options is a right use case for the PIOPS EBS volume?

A. Analytics
B. System boot volume
C. Mongo DB
D. Log processing

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Provisioned IOPS volumes are designed to meet the needs of I/O-intensive workloads, particularly database workloads that are sensitive to storage performance and consistency in random access I/O throughput. Provisioned IOPS volumes are designed to meet the needs of I/O-intensive workloads, particularly database workloads, that are sensitive to storage performance and consistency in random access I/O throughput business applications, database workloads, such as NoSQL DB, RDBMS, etc.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html

QUESTION 211
A company is creating an account strategy so that they can begin using AWS. The Security team will provide each team with the permissions they need to follow the principle or least privileged access. Teams would like to keep their resources isolated from other groups, and the Finance team would like each team's resource usage separated for billing purposes.
Which account creation process meets these requirements and allows for changes?

A. Create a new AWS Organizations account. Create groups in Active Directory and assign them to roles in AWS to grant federated access. Require each team to tag their resources, and separate bills based on tags. Control access to resources through IAM granting the minimally required privilege.
B. Create individual accounts for each team. Assign the security account as the master account, and enable consolidated billing for all other accounts. Create a cross-account role for security to manage accounts, and send logs to a bucket in the security account.
C. Create a new AWS account, and use AWS Service Catalog to provide teams with the required resources. Implement a third-party billing solution to provide the Finance team with the resource use for each team based on tagging. Isolate resources using IAM to avoid account sprawl. Security will control and monitor logs and permissions.
D. Create a master account for billing using Organizations, and create each team's account from that master account.
Create a security account for logs and cross-account access. Apply service control policies on each account, and grant the Security team cross-account access to all accounts. Security will create IAM policies for each account to maintain least privilege access.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
By creating individual IAM users for people accessing your account, you can give each IAM user a unique set of security credentials. You can also grant different permissions to each IAM user. If necessary, you can change or revoke an IAM user's permissions anytime. (If you give out your root user credentials, it can be difficult to revoke them, and it is impossible to restrict their permissions.)
Reference: https://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html

QUESTION 212
Which of the following statements is correct about the number of security groups and rules applicable for an EC2-Classic instance and an EC2-VPC network interface?

A. In EC2-Classic, you can associate an instance with up to 5 security groups and add up to 50 rules to a security group. In EC2-VPC, you can associate a network interface with up to 500 security groups and add up to 100 rules to a security group.
B. In EC2-Classic, you can associate an instance with up to 500 security groups and add up to 50 rules to a security group.
In EC2-VPC, you can associate a network interface with up to 5 security groups and add up to 100 rules to a security group.
C. In EC2-Classic, you can associate an instance with up to 5 security groups and add up to 100 rules to a security group. In EC2-VPC, you can associate a network interface with up to 500 security groups and add up to 50 rules to a security group.
D. In EC2-Classic, you can associate an instance with up to 500 security groups and add up to 100 rules to a security group.
In EC2-VPC, you can associate a network interface with up to 5 security groups and add up to 50 rules to a security group.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A security group acts as a virtual firewall that controls the traffic for one or more instances. When you launch an instance, you associate one or more security groups with the instance. You add rules to each security group that allow traffic to or from its associated instances. If you're using EC2-Classic, you must use security groups created specifically for EC2-Classic. In EC2-Classic, you can associate an instance with up to 500 security groups and add up to 100 rules to a security group. If you're using EC2-VPC, you must use security groups created specifically for your VPC. In EC2-VPC, you can associate a network interface with up to 5 security groups and add up to 50 rules to a security group.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html

QUESTION 213
An organization is setting up RDS for their applications. The organization wants to secure RDS access with VPC.
Which of the following options is not required while designing the RDS with VPC?

A. The organization must create a subnet group with public and private subnets. Both the subnets can be in the same or separate AZ.
B. The organization should keep minimum of one IP address in each subnet reserved for RDS failover.

C. If the organization is connecting RDS from the internet it must enable the VPC attributes DNS hostnames and DNS resolution.
D. The organization must create a subnet group with VPC using more than one subnet which are a part of separate AZs.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources, such as RDS into a virtual network that the user has defined. Subnets are segments of a VPC's IP address range that the user can designate to a group of VPC resources based on security and operational needs. A DB subnet group is a collection of subnets (generally private) that the user can create in a VPC and assign to the RDS DB instances. A DB subnet group allows the user to specify a particular VPC when creating the DB instances.
Each DB subnet group should have subnets in at least two Availability Zones in a given region. If the RDS instance is required to be accessible from the internet the organization must enable the VPC attributes, DNS hostnames and DNS resolution. For each RDS DB instance that the user runs in a VPC, he should reserve at least one address in each subnet in the DB subnet group for use by Amazon RDS for recovery actions.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.html

QUESTION 214
An organization is planning to use NoSQL DB for its scalable data needs. The organization wants to host an application securely in AWS VPC.
What action can be recommended to the organization?

A. The organization should setup their own NoSQL cluster on the AWS instance and configure route tables and subnets.
B. The organization should only use a DynamoDB because by default it is always a part of the default subnet provided by AWS.
C. The organization should use a DynamoDB while creating a table within the public subnet.
D. The organization should use a DynamoDB while creating a table within a private subnet.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Currently VPC does not support DynamoDB. Thus, if the user wants to implement VPC, he has to setup his own NoSQL DB within the VPC.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Introduction.html

QUESTION 215
MapMySite is setting up a web application in the AWS VPC. The organization has decided to use an AWS RDS instead of using its own DB instance for HA and

DR requirements. The organization also wants to secure RDS access. How should the web application be setup with RDS?

A. Create a VPC with one public and one private subnet. Launch an application instance in the public subnet while RDS is launched in the private subnet.
B. Setup a public and two private subnets in different AZs within a VPC and create a subnet group. Launch RDS with that subnet group.
C. Create a network interface and attach two subnets to it. Attach that network interface with RDS while launching a DB instance.
D. Create two separate VPCs and launch a Web app in one VPC and RDS in a separate VPC and connect them with VPC peering.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources, such as RDS into a virtual network that the user has defined. Subnets are segments of a VPC's IP address range that the user can designate to a group of VPC resources based on the security and operational needs.
A DB subnet group is a collection of subnets (generally private) that a user can create in a VPC and assign to the RDS DB instances. A DB subnet group allows the user to specify a particular VPC when creating the DB instances. Each DB subnet group should have subnets in at least two Availability Zones in a given region.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.html

QUESTION 216
A company operates an on-premises software-as-a-service (SaaS) solution that ingests several files daily. The company provides multiple public SFTP endpoints to its customers to facilitate the file transfers. The customers add the SFTP endpoint IP addresses to their firewall allow list for outbound traffic. Changes to the SFTP endpoint IP addresses are not permitted.
The company wants to migrate the SaaS solution to AWS and decrease the operational overhead of the file transfer service.
Which solution meets these requirements?

A. Register the customer-owned block of IP addresses in the company's AWS account. Create Elastic IP addresses from the address pool and assign them to an AWS Transfer for SFTP endpoint. Use AWS Transfer to store the files in Amazon S3.
B. Add a subnet containing the customer-owned block of IP addresses to a VPC. Create Elastic IP addresses from the address pool and assign them to an Application Load Balancer (ALB). Launch EC2 instances hosting FTP services in an Auto Scaling group behind the ALStore the files in attached Amazon Elastic Block Store (Amazon EBS) volumes.
C. Register the customer-owned block of IP addresses with Amazon Route 53. Create alias records in Route 53 that point to a Network Load Balancer (NLB). Launch EC2 instances hosting FTP services in an Auto Scaling group behind the NLB.
Store the files in Amazon S3.
D. Register the customer-owned block of IP addresses in the company's AWS account. Create Elastic IP addresses from the address pool and assign them to an Amazon S3 VPC endpoint. Enable SFTP support on the S3 bucket.

Correct Answer: A

Section: (none)
Explanation

Explanation/Reference:

QUESTION 217
A fitness tracking company serves users around the world, with its primary markets in North America and Asia. The company needs to design an infrastructure for its read-heavy user authorization application with the following requirements:
Be resilient to problem with the application in any Region.
Write to a database in a single Region.
Read from multiple Regions.
Support resiliency across application tiers in each Region.
Support the relational database semantics reflected in the application.
Which combination of steps should a solutions architect take? (Choose two.)

A. Use an Amazon Route 53 geoproximity routing policy combined with a multivalue answer routing policy.
B. Deploy web, application, and MySQL database servers to Amazon EC2 instance in each Region. Set up the application so that reads and writes are local to the Region. Create snapshots of the web, application, and database servers and store the snapshots in an Amazon S3 bucket in both Regions. Set up cross-Region replication for the database layer.
C. Use an Amazon Route 53 geolocation routing policy combined with a failover routing policy.
D. Set up web, application, and Amazon RDS for MySQL instances in each Region. Set up the application so that reads are local and writes are partitioned based on the user. Set up a Multi-AZ failover for the web, application, and database servers.
Set up cross-Region replication for the database layer.
E. Set up active-active web and application servers in each Region. Deploy an Amazon Aurora global database with clusters in each Region. Set up the application to use the in-Region Aurora database endpoints. Create snapshots of the web application servers and store them in an Amazon S3 bucket in both Regions.

Correct Answer: BD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 218
An organization is planning to setup a management network on the AWS VPC. The organization is trying to secure the webserver on a single VPC instance such that it allows the internet traffic as well as the back-end management traffic. The organization wants to make so that the back end management network interface can receive the SSH traffic only from a selected IP range, while the internet facing webserver will have an IP address which can receive traffic from all the internet IPs.

How can the organization achieve this by running web server on a single instance?

A. It is not possible to have two IP addresses for a single instance.
B. The organization should create two network interfaces with the same subnet and security group to assign separate IPs to each network interface.
C. The organization should create two network interfaces with separate subnets so one instance can have two subnets and the respective security groups for controlled access.
D. The organization should launch an instance with two separate subnets using the same network interface which allows to have a separate CIDR as well as security groups.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. An Elastic Network Interface (ENI) is a virtual network interface that the user can attach to an instance in a VPC. The user can create a management network using two separate network interfaces. For the present scenario it is required that the secondary network interface on the instance handles the public facing traffic and the primary network interface handles the back-end management traffic and it is connected to a separate subnet in the VPC that has more restrictive access controls. The public facing interface, which may or may not be behind a load balancer, has an associated security group to allow access to the server from the internet while the private facing interface has an associated security group allowing SSH access only from an allowed range of IP addresses either within the VPC or from the internet, a private subnet within the VPC or a virtual private gateway.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html

QUESTION 219
A company has an application behind a load balancer with enough Amazon EC2 instances to satisfy peak demand. Scripts and third-party deployment solutions are used to configure EC2 instances when demand increases or an instance fails. The team must periodically evaluate the utilization of the instance types to ensure that the correct sizes are deployed.
How can this workload be optimized to meet these requirements?

A. Use CloudFormer to create AWS CloudFormation stacks from the current resources. Deploy that stack by using AWS CloudFormation in the same region. Use Amazon CloudWatch alarms to send notifications about underutilized resources to provide cost-savings suggestions.
B. Create an Auto Scaling group to scale the instances, and use AWS CodeDeploy to perform the configuration. Change from a load balancer to an Application Load Balancer. Purchase a third-party product that provides suggestions for cost savings on AWS resources.
C. Deploy the application by using AWS Elastic Beanstalk with default options. Register for an AWS Support Developer plan.
Review the instance usage for the application by using Amazon CloudWatch, and identify less expensive instances that can handle the load. Hold monthly meetings to review new instance types and determine whether Reserved Instances should be purchased.
D. Deploy the application as a Docker image by using Amazon ECS. Set up Amazon EC2 Auto Scaling and Amazon ECS scaling. Register for AWS Business Support and use Trusted Advisor checks to provide suggestions on cost savings.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 220
A company has multiple AWS accounts and manages these accounts which AWS Organizations. A developer was given IAM user credentials to access AWS resources. The developer should have readonly access to all Amazon S3 buckets in the account. However, when the developer tries to access the S3 buckets from the console, they receive an access denied error message with no bucket listed.
A solution architect reviews the permissions and finds that the developer's IAM user is listed as having read-only access to all S3 buckets in the account. Which additional steps should the solutions architect take to troubleshoot the issue? (Choose two.)

A. Check the bucket policies for all S3 buckets.
B. Check the ACLs for all S3 buckets.
C. Check the SCPs set at the organizational units (OUs).
D. Check for the permissions boundaries set for the IAM user.
E. Check if an appropriate IAM role is attached to the IAM user.

Correct Answer: DE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 221
A company has a photo sharing social networking application. To provide a consistent experience for users, the company performs some image processing on the photos uploaded by users before publishing on the application. The image processing is implemented using a set of Python libraries.
The current architecture is as follows:
The image processing Python code runs in a single Amazon EC2 instance and stores the processed images in an Amazon S3 bucket named ImageBucket. The front-end application, hosted in another bucket, loads the images from ImageBucket to display to users.
With plans for global expansion, the company wants to implement changes in its existing architecture to be able to scale for increased demand on the application and reduce management complexity as the application scales.
Which combination of changes should a solutions architect make? (Choose two.)

A. Place the image processing EC2 instance into an Auto Scaling group.
B. Use AWS Lambda to run the image processing tasks.

C. Use Amazon Rekognition for image processing.
D. Use Amazon CloudFront in front of ImageBucket.
E. Deploy the applications in an Amazon ECS cluster and apply Service Auto Scaling.

Correct Answer: DE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 222
A company wants to analyze log data using date ranges with a custom application running on AWS. The application generates about 10 GB of data every day, which is expected to grow. A Solutions Architect is tasked with storing the data in Amazon S3 and using Amazon Athena to analyze the data.
Which combination of steps will ensure optimal performance as the data grows? (Choose two.)

A. Store each object in Amazon S3 with a random string at the front of each key.
B. Store the data in multiple S3 buckets.
C. Store the data in Amazon S3 in a columnar format, such as Apache Parquet or Apache ORC.
D. Store the data in Amazon S3 in objects that are smaller than 10 MB.
E. Store the data using Apache Hive partitioning in Amazon S3 using a key that includes a date, such as dt=2019-02.

Correct Answer: BC
Section: (none)
Explanation

Explanation/Reference:

QUESTION 223
A company has a VPC with two domain controllers running Active Directory in the default configuration. The VPC DHCP options set is configured to use the IP addresses of the two domain controllers. There is a VPC interface endpoint defined; but instances within the VPC are not able to resolve the private endpoint addresses.
Which strategies would resolve this issue? (Choose two.)

A. Define an outbound Amazon Route 53 Resolver. Set a conditional forward rule for the Active Directory domain to the Active Directory servers. Update the VPC DHCP options set to AmazonProvidedDNS.
B. Update the DNS service on the Active Directory servers to forward all non-authoritative queries to the VPC Resolver.
C. Define an inbound Amazon Route 53 Resolver. Set a conditional forward rule for the Active Directory domain to the Active Directory servers. Update the VPC

DHCP options set to AmazonProvidedDNS.

D. Update the DNS service on the client instances to split DNS queries between the Active Directory servers and the VPC Resolver.

E. Update the DNS service on the Active Directory servers to forward all queries to the VPC Resolver.

Correct Answer: BE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 224
A company runs applications on Amazon EC2 instances. The company plans to begin using an Auto Scaling group for the instances. As part of this transition, a solutions architect must ensure that Amazon CloudWatch Logs automatically collects logs from all new instances. The new Auto Scaling group will use a launch template that includes the Amazon Linux 2 AMI and no key pair.
Which solution meets these requirements?

A. Create an Amazon CloudWatch agent configuration for the workload. Store the CloudWatch agent configuration in an Amazon S3 bucket. Write an EC2 user data script to fetch the configuration file from Amazon S3. Configure the CloudWatch agent on the instance during initial boot.

B. Create an Amazon CloudWatch agent configuration for the workload in AWS Systems Manager Parameter Store. Create a Systems Manager document that installs and configures the CloudWatch agent by using the configuration. Create an Amazon EventBridge (Amazon CloudWatch Events) rule on the default event bus with a Systems Manager Run Command target that runs the document whenever an instance enters the running state.

C. Create an Amazon CloudWatch agent configuration for the workload. Create an AWS Lambda function to install and configure the CloudWatch agent by using AWS Systems Manager Session Manager. Include the agent configuration inside the Lambda package. Create an AWS Config custom rule to identify changes to the EC2 instances and invoke Lambda function.

D. Create an Amazon CloudWatch agent configuration for the workload. Save the CloudWatch agent configuration as part of an AWS Lambda deployment package. Use AWS CloudTrail to capture EC2 tagging events and initiate agent installation.
Use AWS CodeBuild to configure the CloudWatch agent on the instances that run the workload.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/prescriptive-guidance/latest/implementing-logging-monitoring-cloudwatch/installcloudwatch-systems-manager.html

QUESTION 225
A Solutions Architect is designing a highly available and reliable solution for a cluster of Amazon EC2 instances.
The Solutions Architect must ensure that any EC2 instance within the cluster recovers automatically after a system failure.

The solution must ensure that the recovered instance maintains the same IP address.
How can these requirements be met?

A. Create an AWS Lambda script to restart any EC2 instances that shut down unexpectedly.
B. Create an Auto Scaling group for each EC2 instance that has a minimum and maximum size of 1.
C. Create a new t2.micro instance to monitor the cluster instances. Configure the t2.micro instance to issue an aws ec2 reboot-instances command upon failure.
D. Create an Amazon CloudWatch alarm for the StatusCheckFailed_System metric, and then configure an EC2 action to recover the instance.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-recover.html

QUESTION 226
A bank is designing an online customer service portal where customers can chat with customer service agents. The portal is required to maintain a 15-minute RPO or RTO in case of a regional disaster. Banking regulations require that all customer service chat transcripts must be preserved on durable storage for at least 7 years, chat conversations must be encrypted inflight, and transcripts must be encrypted at rest. The Data Loss Prevention team requires that data at rest must be encrypted using a key that the team controls, rotates, and revokes.
Which design meets these requirements?

A. The chat application logs each chat message into Amazon CloudWatch Logs. A scheduled AWS Lambda function invokes a CloudWatch Logs CreateExportTask every 5 minutes to export chat transcripts to Amazon S3. The S3 bucket is configured for cross-region replication to the backup region. Separate AWS KMS keys are specified for the CloudWatch Logs group and the S3 bucket.
B. The chat application logs each chat message into two different Amazon CloudWatch Logs groups in two different regions, with the same AWS KMS key applied. Both CloudWatch Logs groups are configured to export logs into an Amazon Glacier vault with a 7-year vault lock policy with a KMS key specified.
C. The chat application logs each chat message into Amazon CloudWatch Logs. A subscription filter on the CloudWatch Logs group feeds into an Amazon Kinesis Data Firehose which streams the chat messages into an Amazon S3 bucket in the backup region. Separate AWS KMS keys are specified for the CloudWatch Logs group and the Kinesis Data Firehose.
D. The chat application logs each chat message into Amazon CloudWatch Logs. The CloudWatch Logs group is configured to export logs into an Amazon Glacier vault with a 7-year vault lock policy. Glacier cross-region replication mirrors chat archives to the backup region. Separate AWS KMS keys are specified for the CloudWatch Logs group and the Amazon Glacier vault.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/AmazonS3/latest/dev/replication.html

QUESTION 227
How can multiple compute resources be used on the same pipeline in AWS Data Pipeline?

A. You can use multiple compute resources on the same pipeline by defining multiple cluster objects in your definition file and associating the cluster to use for each activity via its runs On field.
B. You can use multiple compute resources on the same pipeline by defining multiple cluster definition files
C. You can use multiple compute resources on the same pipeline by defining multiple clusters for your activity.
D. You cannot use multiple compute resources on the same pipeline.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Multiple compute resources can be used on the same pipeline in AWS Data Pipeline by defining multiple cluster objects in your definition file and associating the cluster to use for each activity via its runs On field, which allows pipelines to combine AWS and on premise resources, or to use a mix of instance types for their activities.
Reference:
https://aws.amazon.com/datapipeline/faqs/

QUESTION 228
A user has created a VPC with CIDR 20.0.0.0/16. The user has created one subnet with CIDR 20.0.0.0/16 in this VPC. The user is trying to create another subnet with the same VPC for CIDR 20.0.0.1/24.
What will happen in this scenario?

A. The VPC will modify the first subnet CIDR automatically to allow the second subnet IP range
B. The second subnet will be created
C. It will throw a CIDR overlaps error
D. It is not possible to create a subnet with the same CIDR as VPC

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. The user can create a subnet with the same size of VPC. However, he cannot create any other subnet since the CIDR of the second subnet will conflict with the first subnet.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

QUESTION 229
A company is using an Amazon EMR cluster to run its big data jobs. The cluster's jobs are invoked by AWS Step Functions Express Workflows that consume various Amazon Simple Queue Service (Amazon SQS) queues. The workload of this solution is variable and unpredictable. Amazon CloudWatch metrics show that the cluster's peak utilization is only 25% at times and that the cluster sits idle the rest of the time.
A solutions architect must optimize the costs of the cluster without negatively impacting the time it takes to run the various jobs.
What is the MOST cost-effective solution that meets these requirements?

A. Modify the EMR cluster by turning on automatic scaling of the core nodes and task nodes with a custom policy that is based on cluster utilization. Purchase Reserved Instance capacity to cover the master node.
B. Modify the EMR cluster to use an instance fleet of Dedicated On-Demand Instances for the master node and core nodes, and to use Spot Instances for the task nodes. Define target capacity for each node type to cover the load.
C. Purchase Reserved Instances for the master node and core nodes. Terminate all existing task nodes in the EMR cluster.
D. Modify the EMR cluster to use capacity-optimized Spot Instances and a diversified task fleet. Define target capacity for each node type with a mix of On-Demand Instances and Spot Instances.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-instance-fleet.html

QUESTION 230
A solutions architect needs to migrate 50 TB of NFS data to Amazon S3. The files are on several NFS file servers on corporate network. These are dense file systems containing tens of millions of small files. The system operators have configured the file interface on an AWS Snowball Edge device and are using a shell script to copy data.
Developers report that copying the data to the Snowball Edge device is very slow. The solutions architect suspects this may be related to the overhead of encrypting all the small files and transporting them over the network.
Which changes can be made to speed up the data transfer?

A. Cluster two Snowball Edge devices together to increase the throughput of the devices.
B. Change the solution to use the S3 Adapter instead of the file interface on the Snowball Edge device.

C. Increase the number of parallel copy jobs to increase the throughput of the Snowball Edge device.
D. Connect directly to the USB interface on the Snowball Edge device and copy the files locally.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 231
To scale out the AWS resources using manual AutoScaling, which of the below mentioned parameters should the user change?

A. Current capacity
B. Desired capacity
C. Preferred capacity
D. Maximum capacity

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The Manual Scaling as part of Auto Scaling allows the user to change the capacity of Auto Scaling group. The user can add / remove EC2 instances on the fly.
To execute manual scaling, the user should modify the desired capacity. AutoScaling will adjust instances as per the requirements.
Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/as-manual-scaling.html

QUESTION 232
Select the correct statement about Amazon ElastiCache.

A. It makes it easy to set up, manage, and scale a distributed in-memory cache environment in the cloud.
B. It allows you to quickly deploy your cache environment only if you install software.
C. It does not integrate with other Amazon Web Services.
D. It cannot run in the Amazon Virtual Private Cloud (Amazon VPC) environment.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
ElastiCache is a web service that makes it easy to set up, manage, and scale a distributed in memory cache environment in the cloud. It provides a high-performance, scalable, and cost- effective caching solution, while removing the complexity associated with deploying and managing a distributed cache environment. With ElastiCache, you can quickly deploy your cache environment, without having to provision hardware or install software.
Reference: http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/WhatIs.html

QUESTION 233
An organization has a write-intensive mobile application that uses Amazon API Gateway, AWS Lambda, and Amazon DynamoDB. The application has scaled well, however, costs have increased exponentially because of higher than anticipated Lambda costs. The application's use is unpredictable, but there has been a steady 20% increase in utilization every month.
While monitoring the current Lambda functions, the Solutions Architect notices that the execution-time averages 4.5 minutes.
Most of the wait time is the result of a high-latency network call to a 3-TB MySQL database server that is on-premises. A VPN is used to connect to the VPC, so the Lambda functions have been configured with a five-minute timeout.
How can the Solutions Architect reduce the cost of the current architecture?

A. Replace the VPN with AWS Direct Connect to reduce the network latency to the on-premises MySQL database.
Enable local caching in the mobile application to reduce the Lambda function invocation calls.
Monitor the Lambda function performance; gradually adjust the timeout and memory properties to lower values while maintaining an acceptable execution time. Offload the frequently accessed records from DynamoDB to Amazon ElastiCache.

B. Replace the VPN with AWS Direct Connect to reduce the network latency to the on-premises MySQL database.
Cache the API Gateway results to Amazon CloudFront.
Use Amazon EC2 Reserved Instances instead of Lambda.
Enable Auto Scaling on EC2, and use Spot Instances during peak times. Enable DynamoDB Auto Scaling to manage target utilization.

C. Migrate the MySQL database server into a Multi-AZ Amazon RDS for MySQL.
Enable caching of the Amazon API Gateway results in Amazon CloudFront to reduce the number of Lambda function invocations.
Monitor the Lambda function performance; gradually adjust the timeout and memory properties to lower values while maintaining an acceptable execution time.
Enable DynamoDB Accelerator for frequently accessed records, and enable the DynamoDB Auto Scaling feature.

D. Migrate the MySQL database server into a Multi-AZ Amazon RDS for MySQL.
Enable API caching on API Gateway to reduce the number of Lambda function invocations.
Continue to monitor the AWS Lambda function performance; gradually adjust the timeout and memory properties to lower values while maintaining an acceptable execution time. Enable Auto Scaling in DynamoDB.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 234
A company currently runs a secure application on Amazon EC2 that takes files from on-premises locations through AWS Direct Connect, processes them, and uploads them to a single Amazon S3 bucket. The application uses HTTPS for encryption in transit to Amazon S3, and S3 server-side encryption to encrypt at rest.
Which of the following changes should the Solutions Architect recommend to make this solution more secure without impeding application's performance?

A. Add a NAT gateway. Update the security groups on the EC2 instance to allow access to and from the S3 IP range only. Configure an S3 bucket policy that allows communication from the NAT gateway's Elastic IP address only.
B. Add a VPC endpoint. Configure endpoint policies on the VPC endpoint to allow access to the required Amazon S3 buckets only. Implement an S3 bucket policy that allows communication from the VPC's source IP range only.
C. Add a NAT gateway. Update the security groups on the EC2 instance to allow access to and from the S3 IP range only. Configure an S3 bucket policy that allows communication from the source public IP address of the on-premises network only.
D. Add a VPC endpoint. Configure endpoint policies on the VPC endpoint to allow access to the required S3 buckets only. Implement an S3 bucket policy that allows communication from the VPC endpoint only.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://docs.aws.amazon.com/AmazonS3/latest/dev/example-bucket-policies-vpc-endpoint.html

QUESTION 235
A company is building a sensor data collection pipeline in which thousands of sensors write data to an Amazon Simple Queue Service (Amazon SQS) queue every minute. The queue is processed by an AWS Lambda function that extracts a standard set of metrics from the sensor data. The company wants to send the data to Amazon CloudWatch. The solution should allow for viewing individual and aggregate sensor metrics and interactively querying the sensor log data using CloudWatch Logs Insights.
What is the MOST cost-effective solution that meets these requirements?

A. Write the processed data to CloudWatch Logs in the CloudWatch embedded metric format.
B. Write the processed data to CloudWatch Logs. Then write the data to CloudWatch by using the PutMetricData API call.
C. Write the processed data to CloudWatch Logs in a structured format. Create a CloudWatch metric filter to parse the logs and publish the metrics to CloudWatch with dimensions to uniquely identify a sensor.
D. Configure the CloudWatch Logs agent for AWS Lambda. Output the metrics for each sensor in statsd format with tags to uniquely identify a sensor. Write the processed data to CloudWatch Logs.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 236
A Solutions Architect is designing a network solution for a company that has applications running in a data center in Northern Virginia. The applications in the company's data center require predictable performance to applications running in a virtual private cloud (VPC) located in us-east-1, and a secondary VPC in us-west-2 within the same account. The company data center is collocated in an AWS Direct Connect facility that serves the us-east-1 region. The company has already ordered an AWS Direct Connect connection and a cross-connect has been established.
Which solution will meet the requirements at the LOWEST cost?

A. Provision a Direct Connect gateway and attach the virtual private gateway (VGW) for the VPC in us-east-1 and the VGW for the VPC in us-west-2. Create a private VIF on the Direct Connect connection and associate it to the Direct Connect gateway.
B. Create private VIFs on the Direct Connect connection for each of the company's VPCs in the us-east-1 and us-west-2 regions. Configure the company's data center router to connect directly with the VPCs in those regions via the private VIFs.
C. Deploy a transit VPC solution using Amazon EC2-based router instances in the us-east-1 region. Establish IPsec VPN tunnels between the transit routers and virtual private gateways (VGWs) located in the us-east-1 and us-west-2 regions, which are attached to the company's VPCs in those regions. Create a public VIF on the Direct Connect connection and establish IPsec VPN tunnels over the public VIF between the transit routers and the company's data center router.
D. Order a second Direct Connect connection to a Direct Connect facility with connectivity to the us-west-2 region. Work with a partner to establish a network extension link over dark fiber from the Direct Connect facility to the company's data center.
Establish private VIFs on the Direct Connect connections for each of the company's VPCs in the respective regions.
Configure the company's data center router to connect directly with the VPCs in those regions via the private VIFs.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://aws.amazon.com/blogs/aws/new-aws-direct-connect-gateway-inter-region-vpc-access/

QUESTION 237
An organization is setting a website on the AWS VPC. The organization has blocked a few IPs to avoid a D-DOS attack.
How can the organization configure that a request from the above mentioned IPs does not access the application instances?

A. Create an IAM policy for VPC which has a condition to disallow traffic from that IP address.

B. Configure a security group at the subnet level which denies traffic from the selected IP.
C. Configure the security group with the EC2 instance which denies access from that IP address.
D. Configure an ACL at the subnet which denies the traffic from that IP address.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. AWS provides two features that the user can use to increase security in VPC: security groups and network ACLs. Security group works at the instance level while ACL works at the subnet level. ACL allows both allow and deny rules. Thus, when the user wants to reject traffic from the selected IPs it is recommended to use ACL with subnets.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_ACLs.html

QUESTION 238
A solutions architect is designing a solution to connect a company's on-premises network with all the company's current and future VPCs on AWS. The company is running VPCs in five different AWS Regions and has at least 15 VPCs in each Region.
The company's AWS usage is constantly increasing and will continue to grow. Additionally, all the VPCs throughout all five Regions must be able to communicate with each other.
The solution must maximize scalability and ease of management.
Which solution meets these requirements?

A. Set up a transit gateway in each Region. Establish a redundant AWS Site-to-Site VPN connection between the onpremises firewalls and the transit gateway in the Region that is closest to the onpremises network. Peer all the transit gateways with each other. Connect all the VPCs to the transit gateway in their Region.
B. Create an AWS CloudFormation template for a redundant AWS Site-to-Site VPN tunnel to the on-premises network. Deploy the CloudFormation template for each VPC. Set up VPC peering between all the VPCs for VPC-to-VPC communication.
C. Set up a transit gateway in each Region. Establish a redundant AWS Site-to-Site VPN connection between the onpremises firewalls and each transit gateway. Route traffic between the different Regions through the company's on-premises firewalls. Connect all the VPCs to the transit gateway in their Region.
D. Create an AWS CloudFormation template for a redundant AWS Site-to-Site VPN tunnel to the on-premises network. Deploy the CloudFormation template for each VPC. Route traffic between the different Regions through the company's onpremises firewalls.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 239
Which statement is NOT true about accessing remote AWS region in the US by your AWS Direct Connect which is located in the US?

A. AWS Direct Connect locations in the United States can access public resources in any US region.
B. You can use a single AWS Direct Connect connection to build multi-region services.
C. Any data transfer out of a remote region is billed at the location of your AWS Direct Connect data transfer rate.
D. To connect to a VPC in a remote region, you can use a virtual private network (VPN) connection over your public virtual interface.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
AWS Direct Connect locations in the United States can access public resources in any US region. You can use a single AWS Direct Connect connection to build multi-region services. To connect to a VPC in a remote region, you can use a virtual private network (VPN) connection over your public virtual interface.
To access public resources in a remote region, you must set up a public virtual interface and establish a border gateway protocol (BGP) session. Then your router learns the routes of the other AWS regions in the US. You can then also establish a VPN connection to your VPC in the remote region.
Any data transfer out of a remote region is billed at the remote region data transfer rate.
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/remote_regions.html

QUESTION 240
You are setting up some EBS volumes for a customer who has requested a setup which includes a RAID (redundant array of inexpensive disks). AWS has some recommendations for RAID setups.
Which RAID setup is not recommended for Amazon EBS?

A. RAID 1 only
B. RAID 5 only
C. RAID 5 and RAID 6
D. RAID 0 only

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
With Amazon EBS, you can use any of the standard RAID configurations that you can use with a traditional bare metal server, as long as that particular RAID

configuration is supported by the operating system for your instance. This is because all RAID is accomplished at the software level. For greater I/O performance than you can achieve with a single volume, RAID 0 can stripe multiple volumes together; for on-instance redundancy, RAID 1 can mirror two volumes together. RAID 5 and RAID 6 are not recommended for Amazon EBS because the parity write operations of these RAID modes consume some of the IOPS available to your volumes.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/raid-config.html

QUESTION 241
A company has an application that runs on Amazon EC2 instances in an Amazon EC2 Auto Scaling group. The company uses AWS CodePipeline to deploy the application. The instances that run in the Auto Scaling group are constantly changing because of scaling events.
When the company deploys new application code versions, the company installs the AWS CodeDeploy agent on any new target EC2 instances and associates the instances with the CodeDeploy deployment group. The application is set to go live within the next 24 hours.
What should a solutions architect recommend to automate the application deployment process with the LEAST amount of operational overhead?

A. Configure Amazon EventBridge (Amazon CloudWatch Events) to invoke an AWS Lambda function when a new EC2 instance is launched into the Auto Scaling group. Code the Lambda function to associate the EC2 instances with the CodeDeploy deployment group.
B. Write a script to suspend Amazon EC2 Auto Scaling operations before the deployment of new code. When the deployment is complete, create a new AMI and configure the Auto Scaling group's launch template to use the new AMI for new launches. Resume Amazon EC2 Auto Scaling operations.
C. Create a new AWS CodeBuild project that creates a new AMI that contains the new code. Configure CodeBuild to update the Auto Scaling group's launch template to the new AMI. Run an Amazon EC2 Auto Scaling instance refresh operation.
D. Create a new AMI that has the CodeDeploy agent installed. Configure the Auto Scaling group's launch template to use the new AMI. Associate the CodeDeploy deployment group with the Auto Scaling group instead of the EC2 instances.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-dg.pdf

QUESTION 242
A company runs an e-commerce platform with front-end and e-commerce tiers. Both tiers run on LAMP stacks with the frontend instances running behind a load balancing appliance that has a virtual offering on AWS. Currently, the Operations team uses SSH to log in to the instances to maintain patches and address other concerns. The platform has recently been the target of multiple attacks, including A DDoS attack.
An SQL injection attack.
Several successful dictionary attacks on SSH accounts on the web servers.
The company wants to improve the security of the e-commerce platform by migrating to AWS. The company's Solutions Architects have decided to use the following approach:
Code review the existing application and fix any SQL injection issues.
Migrate the web application to AWS and leverage the latest AWS Linux AMI to address initial security patching.

Install AWS Systems Manager to manage patching and allow the system administrators to run commands on all instances, as needed.
What additional steps will address all of the identified attack types while providing high availability and minimizing risk?

A. Enable SSH access to the Amazon EC2 instances using a security group that limits access to specific IPs. Migrate onpremises MySQL to Amazon RDS Multi-AZ. Install the third-party load balancer from the AWS Marketplace and migrate the existing rules to the load balancer's AWS instances. Enable AWS Shield Standard for DDoS protection.
B. Disable SSH access to the Amazon EC2 instances. Migrate on-premises MySQL to Amazon RDS Multi-AZ. Leverage an Elastic Load Balancer to spread the load and enable AWS Shield Advanced for protection. Add an Amazon CloudFront distribution in front of the website. Enable AWS WAF on the distribution to manage the rules.
C. Enable SSH access to the Amazon EC2 instances through a bastion host secured by limiting access to specific IP addresses. Migrate on-premises MySQL to a self-managed EC2 instance. Leverage an AWS Elastic Load Balancer to spread the load and enable AWS Shield Standard for DDoS protection. Add an Amazon CloudFront distribution in front of the website.
D. Disable SSH access to the EC2 instances. Migrate on-premises MySQL to Amazon RDS Single-AZ. Leverage an AWS Elastic Load Balancer to spread the load. Add an Amazon CloudFront distribution in front of the website. Enable AWS WAF on the distribution to manage the rules.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 243
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 244
A company is migrating some of its applications to AWS. The company wants to migrate and modernize the applications quickly after it finalizes networking and security strategies. The company has set up an AWS Direct Connection connection in a central network account.
The company expects to have hundreds of AWS accounts and VPCs in the near future. The corporate network must be able to access the resources on AWS seamlessly and also must be able to communicate with all the VPCs. The company also wants to route its cloud resources to the internet through its on-premises data center.
Which combination of steps will meet these requirements? (Choose three.)

A. Create a Direct Connect gateway in the central account. In each of the accounts, create an association proposal by using the Direct Connect gateway and the account ID for every virtual private gateway.
B. Create a Direct Connect gateway and a transit gateway in the central network account. Attach the transit gateway to the Direct Connect gateway by using a transit VIF.
C. Provision an internet gateway. Attach the internet gateway to subnets. Allow internet traffic through the gateway.
D. Share the transit gateway with other accounts. Attach VPCs to the transit gateway.
E. Provision VPC peering as necessary.
F. Provision only private subnets. Open the necessary route on the transit gateway and customer gateway to allow outbound internet traffic from AWS to flow through NAT services that run in the data center.

Correct Answer: BDE
Section: (none)
Explanation

Explanation/Reference:

Explanation:
Reference: https://docs.aws.amazon.com/vpc/latest/tgw/tgw-dcg-attachments.html

QUESTION 245
A company is running a line-of-business (LOB) application on AWS to support its users. The application runs in one VPC, with a backup copy in a second VPC in a different AWS Region for disaster recovery. The company has a single AWS Direct Connect connection between its on-premises network and AWS. The connection terminates at a Direct Connect gateway.
All access to the application must originate from the company's on-premises network and traffic must be encrypted in transit through the use of IPsec. The company is routing traffic through a VPN tunnel over the Direct Connect connection to provide the required encryption.
A business continuity audit determines that the Direct Connect connection represents a potential single point of failure for access to the application. The company needs to remediate this issue as quickly as possible.
Which approach will meet these requirements?

A. Order a second Direct Connect connection to a different Direct Connect location. Terminate the second Direct Connect connection at the same Direct Connect gateway.
B. Configure an AWS Site-to-Site VPN connection over the internet. Terminate the VPN connection at a virtual private gateway in the secondary Region.
C. Create a transit gateway. Attach the VPCs to the transit gateway, and connect the transit gateway to the Direct Connect gateway. Configure an AWS Site-to-Site VPN connection, and terminate it at the transit gateway.
D. Create a transit gateway. Attach the VPCs to the transit gateway, and connect the transit gateway to the Direct Connect gateway. Order a second Direct Connect connection, and terminate it at the transit gateway.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 246
A solutions architect is designing the data storage and retrieval architecture for a new application that a company will be launching soon. The application is designed to ingest millions of small records per minute from devices all around the world.
Each record is less than 4 KB in size and needs to be stored in a durable location where it can be retrieved with low latency.
The data is ephemeral and the company is required to store the data for 120 days only, after which the data can be deleted.
The solutions architect calculates that, during the course of a year, the storage requirements would be about 10-15 TB.
Which storage strategy is the MOST cost-effective and meets the design requirements?

A. Design the application to store each incoming record as a single .csv file in an Amazon S3 bucket to allow for indexed retrieval. Configure a lifecycle policy to delete data older than 120 days.
B. Design the application to store each incoming record in an Amazon DynamoDB table properly configured for the scale.
Configure the DynamoDB Time to Live (TTL) feature to delete records older than 120 days.

C. Design the application to store each incoming record in a single table in an Amazon RDS MySQL database. Run a nightly cron job that executes a query to delete any records older than 120 days.
D. Design the application to batch incoming records before writing them to an Amazon S3 bucket. Update the metadata for the object to contain the list of records in the batch and use the Amazon S3 metadata search feature to retrieve the data.
Configure a lifecycle policy to delete the data after 120 days.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 247
AWS _____supports_____ environments as one of the AWS resource types.

A. Elastic Beanstalk; Elastic Beanstalk application
B. CloudFormation; Elastic Beanstalk application
C. Elastic Beanstalk ; CloudFormation application
D. CloudFormation; CloudFormation application

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
AWS CloudFormation and AWS Elastic Beanstalk services are designed to complement each other. AWS CloudFormation supports Elastic Beanstalk application environments as one of the AWS resource types.
Reference: http://aws.amazon.com/cloudformation/faqs/

QUESTION 248
A company has a media catalog with metadata for each item in the catalog. Different types of metadata are extracted from the media items by an application running on AWS Lambda. Metadata is extracted according to a number of rules with the output stored in an Amazon ElastiCache for Redis cluster. The extraction process is done in batches and takes around 40 minutes to complete.
The update process is triggered manually whenever the metadata extraction rules change.
The company wants to reduce the amount of time it takes to extract metadata from its media catalog. To achieve this, a solutions architect has split the single metadata extraction Lambda function into a Lambda function for each type of metadata.
Which additional steps should the solutions architect take to meet the requirements?

A. Create an AWS Step Functions workflow to run the Lambda functions in parallel. Create another Step Functions workflow that retrieves a list of media items and executes a metadata extraction workflow for each one.
B. Create an AWS Batch compute environment for each Lambda function. Configure an AWS Batch job queue for the compute environment. Create a Lambda function to retrieve a list of media items and write each item to the job queue.
C. Create an AWS Step Functions workflow to run the Lambda functions in parallel. Create a Lambda function to retrieve a list of media items and write each item to an Amazon SQS queue. Configure the SQS queue as an input to the Step Functions workflow.
D. Create a Lambda function to retrieve a list of media items and write each item to an Amazon SQS queue. Subscribe the metadata extraction Lambda functions to the SQS queue with a large batch size.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 249
An EC2 instance that performs source/destination checks by default is launched in a private VPC subnet. All security, NACL, and routing definitions are configured as expected. A custom NAT instance is launched.
Which of the following must be done for the custom NAT instance to work?

A. The source/destination checks should be disabled on the NAT instance.
B. The NAT instance should be launched in public subnet.
C. The NAT instance should be configured with a public IP address.
D. The NAT instance should be configured with an elastic IP address.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Each EC2 instance performs source/destination checks by default. This means that the instance must be the source or destination of any traffic it sends or receives. However, a NAT instance must be able to send and receive traffic when the source or destination is not itself. Therefore, you must disable source/ destination checks on the NAT instance.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_NAT_Instance.html#EIP_DisableSrcDestCheck

QUESTION 250
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 EMR 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 251
A company has developed a custom tool used in its workflow that runs within a Docker container. The company must perform manual steps each time the container code is updated to make the container image available to new workflow executions. The company wants to automate this process to eliminate manual effort and ensure a new container image is generated every time the tool code is updated.
Which combination of actions should a solutions architect take to meet these requirements? (Choose three.)

A. Configure an Amazon ECR repository for the tool. Configure an AWS CodeCommit repository containing code for the tool being deployed to the container image in Amazon ECR.
B. Configure an AWS CodeDeploy application that triggers an application version update that pulls the latest tool container image from Amazon ECR, updates the container with code from the source AWS CodeCommit repository, and pushes the updated container image to Amazon ECR.
C. Configuration an AWS CodeBuild project that pulls the latest tool container image from Amazon ECR, updates the container with code from the source AWS CodeCommit repository, and pushes the updated container image to Amazon ECR.
D. Configure an AWS CodePipeline pipeline that sources the tool code from the AWS CodeCommit repository and initiates an AWS CodeDeploy application update.
E. Configure an Amazon EventBridge rule that triggers on commits to the AWS CodeCommit repository for the tool.
Configure the event to trigger an update to the tool container image in Amazon ECR. Push the updated container image to Amazon ECR.
F. Configure an AWS CodePipeline pipeline that sources the tool code from the AWS CodeCommit repository and initiates an AWS CodeBuild build.

Correct Answer: ACD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 252
A large multinational company runs a timesheet application on AWS that is used by staff across the world. The application runs on Amazon EC2 instances in an Auto Scaling group behind an Elastic Load Balancing (ELB) load balancer, and stores data in an Amazon RDS MySQL Multi-AZ database instance.
The CFO is concerned about the impact on the business if the application is not available. The application must not be down for more than two hours, but he solution must be as cost-effective as possible.
How should the Solutions Architect meet the CFO's requirements while minimizing data loss?

A. In another region, configure a read replica and create a copy of the infrastructure. When an issue occurs, promote the read replica and configure as an Amazon RDS Multi-AZ database instance. Update the DNS record to point to the other region's ELB.
B. Configure a 1-day window of 60-minute snapshots of the Amazon RDS Multi-AZ database instance. Create an AWS CloudFormation template of the application infrastructure that uses the latest snapshot. When an issue occurs, use the AWS CloudFormation template to create the environment in another region. Update the DNS record to point to the other region's ELB.
C. Configure a 1-day window of 60-minute snapshots of the Amazon RDS Multi-AZ database instance which is copied to another region. Create an AWS CloudFormation template of the application infrastructure that uses the latest copied snapshot. When an issue occurs, use the AWS CloudFormation template to create the environment in another region.
Update the DNS record to point to the other region's ELB.
D. Configure a read replica in another region. Create an AWS CloudFormation template of the application infrastructure.
When an issue occurs, promote the read replica and configure as an Amazon RDS Multi-AZ database instance and use the AWS CloudFormation template to create the environment in another region using the promoted Amazon RDS instance.
Update the DNS record to point to the other region's ELB.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 253
You are designing an SSL/TLS solution that requires HTTPS clients to be authenticated by the Web server using clientcertificate authentication. The solution must be resilient. Which of the following options would you consider for configuringthe web server infrastructure? (Choose two.)

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:
Explanation:
TCP/443 or HTTPS listener either way you can configure, but you can only upload ssl certificate on HTTPS listener.

QUESTION 254
A company has a large number of AWS accounts in an organization in AWS Organizations. A different business group owns each account. All the AWS accounts are bound by legal compliance requirements that restrict all operations outside the euwest- 2 Region.
The company's security team has mandated the use of AWS Systems Manager Session Manager across all AWS accounts.
Which solution should a solutions architect recommend to meet these requirements?

A. Create an SCP that denies access to all requests that do not target eu-west-2. Use the NotAction element to exempt global services from the restriction. In AWS Organizations, apply the SCP to the root of the organization.
B. Create an SCP that denies access to all requests that do not target eu-west-2. Use the NotAction element to exempt global services from the restriction. For each AWS account, use the AmNotLike condition key to add the ARN of the IAM role that is associated with the Session Manager instance profile to the condition element of the SCP. In AWS Organizations apply, the SCP to the root of the organization.
C. Create an SCP that denies access to all requests that do not target eu-west-2. Use the NotAction element to exempt global services from the restriction. In AWS Organizations, apply the SCP to the root of the organization. In each AWS account, create an IAM permissions boundary that allows access to the IAM role that is associated with the Session Manager instance profile.
D. For each AWS account, create an IAM permissions boundary that denies access to all requests that do not target euwest- 2. For each AWS account, apply the permissions boundary to the IAM role that is associated with the Session Manager instance profile.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_examples_aws_deny-requestedregion.html

QUESTION 255
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 256
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 base 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 257
A company developed a Java application and deployed it to an Apache Tomcat server that runs on Amazon EC2 instances.
The company's Engineering team has implemented AWS CloudFormation and Chef Automate to automate the provisioning of and updates to the infrastructure and configuration of the application in the development, test, and production environments. These implementations have led to significantly improves reliability in releasing changes. The Engineering team reports there are frequent service disruptions due to unexpected errors when updating the application of the Apache Tomcat server.
Which solution will increase the reliability of all releases?

A. Implement a blue/green deployment methodology.
B. Implement the canary release methodology.
C. Configure Amazon CloudFront to serve all requests from the cache while deploying the updates.
D. Implement the all at once deployment methodology.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://medium.com/@tom.tikkle/blue-green-deployments-increasing-safety-reliability-speed-98a5c6b222b0

QUESTION 258
A user has created a VPC with public and private subnets using the VPC Wizard. The VPC has CIDR 20.0.0.0/16. The private subnet uses CIDR 20.0.0.0/24.
Which of the below mentioned entries are required in the main route table to allow the instances in VPC to communicate with each other?

A. Destination : 20.0.0.0/0 and Target : ALL
B. Destination : 20.0.0.0/16 and Target : Local
C. Destination : 20.0.0.0/24 and Target : Local
D. Destination : 20.0.0.0/16 and Target : ALL

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A user can create a subnet with VPC and launch instances inside that subnet. If the user has created a public private subnet, the instances in the public subnet can receive inbound traffic directly from the Internet, whereas the instances in the private subnet cannot. If these subnets are created with Wizard, AWS will create two route tables and attach to the subnets. The main route table will have the entry "Destination: 20.0.0.0/16 and Target: Local", which allows all instances in the VPC to communicate with each other.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario2.html

QUESTION 259
A company in the United States (US) has acquired a company in Europe. Both companies use the AWS Cloud. The US company has built a new application with a microservices architecture. The US company is hosting the application across five VPCs in the us-east-2 Region. The application must be able to access resources in one VPC in the eu-west-1 Region.
However, the application must not be able to access any other VPCs.
The VPCs in both Regions have no overlapping CIDR ranges. All Accounts are already consolidated in one organization in AWS Organizations.
Which solution will meet these requirements MOST cost-effectively?

A. Create one transit gateway in eu-west-1. Attach the VPCs in us-east-2 and the VPC in eu-west-1 to the transit gateway. Create the necessary route entries in each VPC so that the traffic is routed through the transit gateway.
B. Create one transit gateway in each Region. Attach the involved subnets to the regional transit gateway. Create the necessary route entries in the associated route tables for each subnet so that the traffic is routed through the regional transit gateway. Peer the two transit gateways.
C. Create a full mesh VPC peering connection configuration between all the VPCs. Create the necessary route entries in each VPC so that the traffic is routed through the VPC peering connection.
D. Create one VPC peering connection for each VPC in us-east-2 to the VPC in eu-west-1. Create the necessary route entries in each VPC so that the traffic is routed through the VPC peering connection.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/vpc/latest/tgw/how-transit-gateways-work.html

QUESTION 260
Can a user configure a custom health check with Auto Scaling?

A. Yes, but the configured data will not be saved to Auto Scaling.
B. No, only an ELB health check can be configured with Auto Scaling.

C. Yes
D. No

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Auto Scaling can determine the health status of an instance using custom health checks. If you have custom health checks, you can send the information from your health checks to Auto Scaling so that Auto Scaling can use this information. For example, if you determine that an instance is not functioning as expected, you can set the health status of the instance to Unhealthy. The next time that Auto Scaling performs a health check on the instance, it will determine that the instance is unhealthy and then launch a replacement instance.
Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/healthcheck.html

QUESTION 261
A company has decided to move some workloads onto AWS to create a grid environment to run market analytics. The grid will consist of many similar instances, spun-up by a job-scheduling function. Each time a large analytics workload is completed, a new VPC is deployed along with job scheduler and grid nodes. Multiple grids could be running in parallel.
Key requirements are:
Grid instances must communicate with Amazon S3 to retrieve data to be processed.
Grid instances must communicate with Amazon DynamoDB to track intermediate data.
The job scheduler needs only to communicate with the Amazon EC2 API to start new grid nodes.
A key requirement is that the environment has no access to the internet, either directly or via the on-premises proxy.
However, the application needs to be able to seamlessly communicate to Amazon S3, Amazon DynamoDB, and Amazon EC2 API, without the need for reconfiguration for each new deployment. Which of the following should the Solutions Architect do to achieve this target architecture? (Choose three.)

A. Enable VPC endpoints for Amazon S3 and DynamoDB.
B. Disable Private DNS Name Support.
C. Configure the application on the grid instances to use the private DNS name of the Amazon S3 endpoint.
D. Populate the on-premises DNS server with the private IP addresses of the EC2 endpoint.
E. Enable an interface VPC endpoint for EC2.
F. Configure Amazon S3 endpoint policy to permit access only from the grid nodes.

Correct Answer: ACE
Section: (none)
Explanation

Explanation/Reference:

Explanation:
Reference: https://aws.amazon.com/premiumsupport/knowledge-center/connect-s3-vpc-endpoint/
https://docs.aws.amazon.com/vpc/latest/userguide/vpce-interface.html

QUESTION 262
With Amazon Elastic MapReduce (Amazon EMR) you can analyze and process vast amounts of data. The cluster is managed using an open-source framework called Hadoop. You have set up an application to run Hadoop jobs. The application reads data from DynamoDB and generates a temporary file of 100 TBs. The whole process runs for 30 minutes and the output of the job is stored to S3.
Which of the below mentioned options is the most cost effective solution in this case?

A. Use Spot Instances to run Hadoop jobs and configure them with EBS volumes for persistent data storage.
B. Use Spot Instances to run Hadoop jobs and configure them with ethereal storage for output file storage.
C. Use an on demand instance to run Hadoop jobs and configure them with EBS volumes for persistent storage.
D. Use an on demand instance to run Hadoop jobs and configure them with ephemeral storage for output file storage.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
AWS EC2 Spot Instances allow the user to quote his own price for the EC2 computing capacity. The user can simply bid on the spare Amazon EC2 instances and run them whenever his bid exceeds the current Spot Price. The Spot Instance pricing model complements the On-Demand and Reserved Instance pricing models, providing potentially the most cost-effective option for obtaining compute capacity, depending on the application. The only challenge with a Spot Instance is data persistence as the instance can be terminated whenever the spot price exceeds the bid price. In the current scenario a Hadoop job is a temporary job and does not run for a longer period. It fetches data from a persistent DynamoDB. Thus, even if the instance gets terminated there will be no data loss and the job can be re-run. As the output files are large temporary files, it will be useful to store data on ethereal storage for cost savings.
Reference: http://aws.amazon.com/ec2/purchasing-options/spot-instances/

QUESTION 263
A company has a legacy application running on servers on premises. To increase the application's reliability, the company wants to gain actionable insights using application logs. A Solutions Architect has been given following requirements for the solution:
Aggregate logs using AWS.
Automate log analysis for errors.
Notify the Operations team when errors go beyond a specified threshold.
What solution meets the requirements?

A. Install Amazon Kinesis Agent on servers, send logs to Amazon Kinesis Data Streams and use Amazon Kinesis Data Analytics to identify errors, create an Amazon CloudWatch alarm to notify the Operations team of errors
B. Install an AWS X-Ray agent on servers, send logs to AWS Lambda and analyze them to identify errors, use Amazon CloudWatch Events to notify the

Operations team of errors.

C. Install Logstash on servers, send logs to Amazon S3 and use Amazon Athena to identify errors, use sendmail to notify the Operations team of errors.

D. Install the Amazon CloudWatch agent on servers, send logs to Amazon CloudWatch Logs and use metric filters to identify errors, create a CloudWatch alarm to notify the Operations team of errors.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html
https://docs.aws.amazon.com/kinesis-agent-windows/latest/userguide/what-is-kinesis-agent-windows.html

QUESTION 264
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, Soft state, 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 onpremises 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 265
A user has created a VPC with public and private subnets. The VPC has CIDR 20.0.0.0/16. The private subnet uses CIDR 20.0.1.0/24 and the public subnet uses CIDR 20.0.0.0/24. The user is planning to host a web server in the public subnet (port 80) and a DB server in the private subnet (port 3306). The user is configuring a security group of the NAT instance.
Which of the below mentioned entries is not required in NAT's security group for the database servers to connect to the Internet for software updates?

A. For Outbound allow Destination: 0.0.0.0/0 on port 443
B. For Inbound allow Source: 20.0.1.0/24 on port 80
C. For Inbound allow Source: 20.0.0.0/24 on port 80
D. For Outbound allow Destination: 0.0.0.0/0 on port 80

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A user can create a subnet with VPC and launch instances inside that subnet. If the user has created a public private subnet to host the web server and DB server respectively, the user should configure that the instances in the private subnet can connect to the internet using the NAT instances. The user should first configure that NAT can receive traffic on ports 80 and 443 from the private subnet. Thus, allow ports 80 and 443 in Inbound for the private subnet 20.0.1.0/24. Now to route this traffic to the internet configure ports 80 and 443 in Outbound with destination 0.0.0.0/0. The NAT should not have an entry for the public subnet CIDR.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario2.html

QUESTION 266
A video processing company wants to build a machine learning (ML) model by using 600 TB of compressed data that is stored as thousands of files in the company's on-premises network attached storage system. The company does not have the necessary compute resources on premises for ML experiments and wants to use AWS.
The company needs to complete the data transfer to AWS within 3 weeks. The data transfer will be a one-time transfer. The data must be encrypted in transit. The measured upload speed of the company's internet connection is 100 Mbps, and multiple departments share the connection.
Which solution will meet these requirements MOST cost-effectively?

A. Order several AWS Snowball Edge Storage Optimized devices by using the AWS Management Console. Configure the devices with a destination S3 bucket. Copy the data to the devices. Ship the devices back to AWS.
B. Set up a 10 Gbps AWS Direct Connect connection between the company location and the nearest AWS Region. Transfer the data over a VPN connection into the Region to store the data in Amazon S3.
C. Create a VPN connection between the on-premises network storage and the nearest AWS Region. Transfer the data over the VPN connection.
D. Deploy an AWS Storage Gateway file gateway on premises. Configure the file gateway with a destination S3 bucket. Copy the data to the file gateway.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

Explanation:
Reference: https://aws.amazon.com/blogs/networking-and-content-delivery/building-a-global-network-using-aws-transitgateway-inter-region-peering/

QUESTION 267
A company runs its containerized batch jobs on Amazon ECS. The jobs are scheduled by submitting a container image, a task definition, and the relevant data to an Amazon S3 bucket. Container images may be unique per job. Running the jobs as quickly as possible is of utmost importance, so submitting job artifacts to the S3 bucket triggers the job to run immediately. Sometimes there may be no jobs running at all. However, jobs of any size can be submitted with no prior warning to the IT Operations team. Job definitions include CPU and memory resource requirements.
What solution will allow the batch jobs to complete as quickly as possible after being scheduled?

A. Schedule the jobs on an Amazon ECS cluster using the Amazon EC2 launch type. Use Service Auto Scaling to increase or decrease the number of running tasks to suit the number of running jobs.
B. Schedule the jobs directly on EC2 instances. Use Reserved Instances for the baseline minimum load, and use On- Demand Instances in an Auto Scaling group to scale up the platform based on demand.
C. Schedule the jobs on an Amazon ECS cluster using the Fargate launch type. Use Service Auto Scaling to increase or decrease the number of running tasks to suit the number of running jobs.
D. Schedule the jobs on an Amazon ECS cluster using the Fargate launch type. Use Spot Instances in an Auto Scaling group to scale the platform based on demand. Use Service Auto Scaling to increase or decrease the number of running tasks to suit the number of running jobs.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 268
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 269
A user has created an AWS AMI. The user wants the AMI to be available only to his friend and not anyone else. How can the user manage this?

A. Share the AMI with the community and setup the approval workflow before anyone launches it.
B. It is not possible to share the AMI with the selected user.
C. Share the AMI with a friend's AWS account ID.
D. Share the AMI with a friend's AWS login ID.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In Amazon Web Services, if a user has created an AMI and wants to share with his friends and colleagues he can share the AMI with their AWS account ID.
Once the AMI is shared the other user can access it from the community AMIs under private AMIs options.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/sharingamis-explicit.html

QUESTION 270
A customer has a website which shows all the deals available across the market. The site experiences a load of 5 large EC2 instances generally.
However, a week before Thanksgiving vacation they encounter a load of almost 20 large instances. The load during that period varies over the day based on the office timings.
Which of the below mentioned solutions is cost effective as well as help the website achieve better performance?

A. Setup to run 10 instances during the pre-vacation period and only scale up during the office time by launching 10 more instances using the AutoScaling schedule.
B. Keep only 10 instances running and manually launch 10 instances every day during office hours.
C. During the pre-vacation period setup 20 instances to run continuously.
D. During the pre-vacation period setup a scenario where the organization has 15 instances running and 5 instances to scale up and down using Auto Scaling based on the network I/O policy.

Correct Answer: B

Section: (none)
Explanation

Explanation/Reference:
Explanation:
AWS provides an on demand, scalable infrastructure. AWS EC2 allows the user to launch On- Demand instances and the organization should create an AMI of the running instance. When the organization is experiencing varying loads and the time of the load is not known but it is higher than the routine traffic it is recommended that the organization launches a few instances beforehand and then setups AutoScaling with policies which scale up and down as per the EC2 metrics, such as Network I/O or CPU utilization. If the organization keeps all 10 additional instances as a part of the AutoScaling policy sometimes during a sudden higher load it may take time to launch instances and may not give an optimal performance. This is the reason it is recommended that the organization keeps an additional 5 instances running and the next 5 instances scheduled as per the AutoScaling policy for cost effectiveness.

QUESTION 271
Complete this statement: "When you load your table directly from an Amazon_____ table, you have the option to control the amount of provisioned throughput you consume."

A. RDS
B. DataPipeline
C. DynamoDB
D. S3

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
When you load your table directly from an Amazon DynamoDB table, you have the option to control the amount of Amazon DynamoDB provisioned throughput you consume.
Reference: http://docs.aws.amazon.com/redshift/latest/dg/t_Loading_tables_with_the_COPY_command.html

QUESTION 272
A company is creating a sequel for a popular online game. A large number of users from all over the world will play the game within the first week after launch. Currently, the game consists of the following components deployed in a single AWS Region:
Amazon S3 bucket that stores game assets
Amazon DynamoDB table that stores player scores
A solutions architect needs to design a multi-Region solution that will reduce latency, improve reliability, and require the least effort to implement.
What should the solutions architect do to meet these requirements?

A. Create an Amazon CloudFront distribution to serve assets from the S3 bucket. Configure S3 Cross-Region Replication.

Create a new DynamoDB table in a new Region. Use the new table as a replica target for DynamoDB global tables.

B. Create an Amazon CloudFront distribution to serve assets from the S3 bucket. Configure S3 Same-Region Replication. Create a new DynamoDB table in a new Region. Configure asynchronous replication between the DynamoDB tables by using AWS Database Migration Service (AWS DMS) with change data capture (CDC).

C. Create another S3 bucket in a new Region, and configure S3 Cross-Region Replication between the buckets. Create an Amazon CloudFront distribution and configure origin failover with two origins accessing the S3 buckets in each Region. Configure DynamoDB global tables by enabling Amazon DynamoDB Streams, and add a replica table in a new Region.

D. Create another S3 bucket in the sine Region, and configure S3 Same-Region Replication between the buckets. Create an Amazon CloudFront distribution and configure origin failover with two origins accessing the S3 buckets. Create a new DynamoDB table in a new Region. Use the new table as a replica target for DynamoDB global tables.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/blogs/publicsector/how-to-meet-business-data-resiliency-s3-cross-region-replication/

QUESTION 273
A company wants to migrate its on-premises data center to the AWS Cloud. This includes thousands of virtualized Linux and Microsoft Windows servers, SAN storage, Java and PHP applications with MYSQL, and Oracle databases. There are many department services hosted either in the same data center or externally. The technical documentation is incomplete and outdated. A solutions architect needs to understand the current environment and estimate the cloud resource costs after the migration. Which tools or services should solutions architect use to plan the cloud migration (Choose three.)

A. AWS Application Discovery Service
B. AWS SMS
C. AWS x-Ray
D. AWS Cloud Adoption Readness Tool (CART)
E. Amazon Inspector
F. AWS Migration Hub

Correct Answer: BCF
Section: (none)
Explanation

Explanation/Reference:

QUESTION 274
A company is adding a new approved external vendor that only supports IPv6 connectivity. The company's backend systems sit in the private subnet of an Amazon VPC. The company uses a NAT gateway to allow these systems to communicate with external vendors over IPv4. Company policy requires systems that communicate with external vendors to use a security group that limits access to only approved external vendors. The virtual private cloud (VPC) uses the default network ACL.
The Systems Operator successfully assigns IPv6 addresses to each of the backend systems. The Systems Operator also updates the outbound security group to include the IPv6 CIDR of the external vendor (destination). The systems within the VPC are able to ping one another successfully over IPv6. However, these systems are unable to communicate with the external vendor.
What changes are required to enable communication with the external vendor?

A. Create an IPv6 NAT instance. Add a route for destination 0.0.0.0/0 pointing to the NAT instance.
B. Enable IPv6 on the NAT gateway. Add a route for destination ::/0 pointing to the NAT gateway.
C. Enable IPv6 on the internet gateway. Add a route for destination 0.0.0.0/0 pointing to the IGW.
D. Create an egress-only internet gateway. Add a route for destination ::/0 pointing to the gateway.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://docs.aws.amazon.com/vpc/latest/userguide/egress-only-internet-gateway.html

QUESTION 275
A company is deploying a public-facing global application on AWS using Amazon CloudFront. The application communicates with an external system. A solutions architect needs to ensure the data is secured during end-to-end transit and at rest.
Which combination of steps will satisfy these requirements? (Choose three.)

A. Create a public certificate for the required domain in AWS Certificate Manager and deploy it to CloudFront, an Application Load Balancer, and Amazon EC2 instances.
B. Acquire a public certificate from a third-party vendor and deploy it to CloudFront, an Application Load Balancer, and Amazon EC2 instances.
C. Provision Amazon EBS encrypted volumes using AWS KMS and ensure explicit encryption of data when writing to Amazon EBS.
D. Provision Amazon EBS encrypted volumes using AWS KMS.
E. Use SSL or encrypt data while communicating with the external system using a VPN.
F. Communicate with the external system using plaintext and use the VPN to encrypt the data in transit.

Correct Answer: ACE
Section: (none)

Explanation

Explanation/Reference:

QUESTION 276
A company has an on-premises Microsoft SQL Server database that writes a nightly 200 GB export to a local drive. The company wants to move the backups to more robust cloud storage on Amazon S3. The company has set up a 10 Gbps AWS Direct Connect connection between the on-premises data center and AWS. Which solution meets these requirements MOST cost-effectively?

A. Create a new S3 bucket. Deploy an AWS Storage Gateway file gateway within the VPC that is connected to the Direct Connect connection. Create a new SMB file share. Write nightly database exports to the new SMB file share.
B. Create an Amazon FSx for Windows File Server Single-AZ file system within the VPC that is connected to the Direct Connect connection. Create a new SMB file share. Write nightly database exports to an SMB file share on the Amazon FSx file system. Enable nightly backups.
C. Create an Amazon FSx for Windows File Server Multi-AZ file system within the VPC that is connected to the Direct Connect connection. Create a new SMB file share. Write nightly database exports to an SMB file share on the Amazon FSx file system. Enable nightly backups.
D. Create a new S3 bucket. Deploy an AWS Storage Gateway volume gateway within the VPC that is connected to the Direct Connect connection. Create a new SMB file share. Write nightly database exports to the new SMB file share on the volume gateway, and automate copies of this data to an S3 bucket.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/blogs/storage/accessing-smb-file-shares-remotely-with-amazon-fsx-for-windows-fileserver/

QUESTION 277
A life sciences company is using a combination of open source tools to manage data analysis workflows and Docker containers running on servers in its on-premises data center to process genomics data. Sequencing data is generated and stored on a local storage area network (SAN), and then the data is processed. The research and development teams are running into capacity issues and have decided to re-architect their genomics analysis platform on AWS to scale based on workload demands and reduce the turnaround time from weeks to days.
The company has a high-speed AWS Direct Connect connection. Sequencers will generate around 200 GB of data for each genome, and individual jobs can take several hours to process the data with ideal compute capacity. The end result will be stored in Amazon S3. The company is expecting 10-15 job requests each day.
Which solution meets these requirements?

A. Use regularly scheduled AWS Snowball Edge devices to transfer the sequencing data into AWS. When AWS receives the Snowball Edge device and the data is loaded into Amazon S3, use S3 events to trigger an AWS Lambda function to process the data.
B. Use AWS Data Pipeline to transfer the sequencing data to Amazon S3. Use S3 events to trigger an Amazon EC2 Auto Scaling group to launch custom-AMI

EC2 instances running the Docker containers to process the data.

C. Use AWS DataSync to transfer the sequencing data to Amazon S3. Use S3 events to trigger an AWS Lambda function that starts an AWS Step Functions workflow. Store the Docker images in Amazon Elastic Container Registry (Amazon ECR) and trigger AWS Batch to run the container and process the sequencing data.

D. Use an AWS Storage Gateway file gateway to transfer the sequencing data to Amazon S3. Use S3 events to trigger an AWS Batch job that executes on Amazon EC2 instances running the Docker containers to process the data.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 278
AWS CloudFormation _____ are special actions you use in your template to assign values to properties that are not available until runtime.

A. intrinsic functions
B. properties declarations
C. output functions
D. conditions declarations

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
AWS CloudFormation intrinsic functions are special actions you use in your template to assign values to properties not available until runtime. Each function is declared with a name enclosed in quotation marks (""), a single colon, and its parameters.
Reference: http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-fuctions-structure.html

QUESTION 279
A hedge fund company is developing a new web application to handle trades. Traders around the world will use the application. The application will handle hundreds of thousands of transactions, especially during overlapping work hours between Europe and the United States.
According to the company's disaster recovery plan, the data that is generated must be replicated to a second AWS Region.
Each transaction item will be less than 100 KB in size. The company wants to simplify the CI/CD pipeline as much as possible.
Which combination of steps will meet these requirements MOST cost-effectively? (Choose two.)

A. Deploy the application in multiple Regions. Use Amazon Route 53 latency-based routing to route users to the nearest deployment.
B. Provision an Amazon Aurora global database to persist data. Use Amazon ElastiCache to improve response time.
C. Provision an Amazon CloudFront domain with the website as an origin. Restrict access to geographies where the usage is expected.
D. Provision an Amazon DynamoDB global table. Use DynamoDB Accelerator (DAX) to improve response time.
E. Provision an Amazon Aurora multi-master cluster to persist data. Use Amazon ElastiCache to improve response time.

Correct Answer: AB
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html

QUESTION 280
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 281
A company collects a steady stream of 10 million data records from 100,000 sources each day. These records are written to an Amazon RDS MySQL DB. A query must produce the daily average of a data source over the past 30 days. There are twice as many reads as writes. Queries to the collected data are for one source ID at a time.
How can the Solutions Architect improve the reliability and cost effectiveness of this solution?

A. Use Amazon Aurora with MySQL in a Multi-AZ mode. Use four additional read replicas.
B. Use Amazon DynamoDB with the source ID as the partition key and the timestamp as the sort key. Use a Time to Live (TTL) to delete data after 30 days.
C. Use Amazon DynamoDB with the source ID as the partition key. Use a different table each day.

D. Ingest data into Amazon Kinesis using a retention period of 30 days. Use AWS Lambda to write data records to Amazon ElastiCache for read access.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html

QUESTION 282
Identify a true statement about using an IAM role to grant permissions to applications running on Amazon EC2 instances.

A. When AWS credentials are rotated; developers have to update only the root Amazon EC2 instance that uses their credentials.
B. When AWS credentials are rotated, developers have to update only the Amazon EC2 instance on which the password policy was applied and which uses their credentials.
C. When AWS credentials are rotated, you don't have to manage credentials and you don't have to worry about long-term security risks.
D. When AWS credentials are rotated, you must manage credentials and you should consider precautions for long-term security risks.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Using IAM roles to grant permissions to applications that run on EC2 instances requires a bit of extra configuration. Because role credentials are temporary and rotated automatically, you don't have to manage credentials, and you don't have to worry about long-term security risks.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/role-usecase-ec2app.html

QUESTION 283
A company is planning to migrate an application from on-premises to AWS. The application currently uses an Oracle database and the company can tolerate a brief downtime of 1 hour when performing the switch to the new infrastructure. As part of the migration, the database engine will be changed to MySQL. A Solutions Architect needs to determine which AWS services can be used to perform the migration while minimizing the amount of work and time required. Which of the following will meet the requirements?

A. Use AWS SCT to generate the schema scripts and apply them on the target prior to migration. Use AWS DMS to analyze the current schema and provide a recommendation for the optimal database engine. Then, use AWS DMS to migrate to the recommended engine. Use AWS SCT to identify what embedded SQL code in the application can be converted and what has to be done manually.
B. Use AWS SCT to generate the schema scripts and apply them on the target prior to migration. Use AWS DMS to begin moving data from the on-premises database to AWS. After the initial copy, continue to use AWS DMS to keep the databases in sync until cutting over to the new database. Use AWS SCT to

identify what embedded SQL code in the application can be converted and what has to be done manually.

C. Use AWS DMS to help identify the best target deployment between installing the database engine on Amazon EC2 directly or moving to Amazon RDS. Then, use AWS DMS to migrate to the platform. Use AWS Application Discovery Service to identify what embedded SQL code in the application can be converted and what has to be done manually.

D. Use AWS DMS to begin moving data from the on-premises database to AWS. After the initial copy, continue to use AWS DMS to keep the databases in sync until cutting over to the new database. Use AWS Application Discovery Service to identify what embedded SQL code in the application can be converted and what has to be done manually.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 284
A user has launched a dedicated EBS backed instance with EC2. You are curious where the EBS volume for this instance will be created.
Which statement is correct about the EBS volume's creation?

A. The EBS volume will not be created on the same tenant hardware assigned to the dedicated instance
B. AWS does not allow a dedicated EBS backed instance launch
C. The EBS volume will be created on the same tenant hardware assigned to the dedicated instance
D. The user can specify where the EBS will be created

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The dedicated instances are Amazon EC2 instances that run in a Virtual Private Cloud (VPC) on hardware that is dedicated to a single customer. When a user launches an Amazon EBS-backed dedicated instance, the EBS volume does not run on single-tenant hardware.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/dedicated-instance.html

QUESTION 285
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 286
A company has a new security policy. The policy requires the company to log any event that retrieves data from Amazon S3 buckets. The company must save these audit logs in a dedicated S3 bucket.
The company created the audit logs S3 bucket in an AWS account that is designated for centralized logging. The S3 bucket has a bucket policy that allows write-only cross-account access.
A solutions architect must ensure that all S3 object-level access is being logged for current S3 buckets and future S3 buckets.
Which solution will meet these requirements?

A. Enable server access logging for all current S3 buckets. Use the audit logs S3 bucket as a destination for audit logs.
B. Enable replication between all current S3 buckets and the audit logs S3 bucket. Enable S3 Versioning in the audit logs S3 bucket.
C. Configure S3 Event Notifications for all current S3 buckets to invoke an AWS Lambda function every time objects are accessed. Store Lambda logs in the audit logs S3 bucket.
D. Enable AWS CloudTrail, and use the audit logs S3 bucket to store logs. Enable data event logging for S3 event sources, current S3 buckets, and future S3 buckets.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/awscloudtrail/latest/userguide/best-practices-security.html

QUESTION 287
An organization is setting up a web application with the JEE stack. The application uses the JBoss app server and MySQL DB. The application has a logging module which logs all the activities whenever a business function of the JEE application is called. The logging activity takes some time due to the large size of

the log file.
If the application wants to setup a scalable infrastructure which of the below mentioned options will help achieve this setup?

A. Host the log files on EBS with PIOPS which will have higher I/O.
B. Host logging and the app server on separate servers such that they are both in the same zone.
C. Host logging and the app server on the same instance so that the network latency will be shorter.
D. Create a separate module for logging and using SQS compartmentalize the module such that all calls to logging are asynchronous.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The organization can always launch multiple EC2 instances in the same region across multiple AZs for HA and DR. The AWS architecture practice recommends compartmentalizing the functionality such that they can both run in parallel without affecting the performance of the main application. In this scenario logging takes a longer time due to the large size of the log file. Thus, it is recommended that the organization should separate them out and make separate modules and make asynchronous calls among them. This way the application can scale as per the requirement and the performance will not bear the impact of logging.
Reference: http://www.awsarchitectureblog.com/2014/03/aws-and-compartmentalization.html

QUESTION 288
In Amazon Cognito what is a silent push notification?

A. It is a push message that is received by your application on a user's device that will not be seen by the user.
B. It is a push message that is received by your application on a user's device that will return the user's geolocation.
C. It is a push message that is received by your application on a user's device that will not be heard by the user.
D. It is a push message that is received by your application on a user's device that will return the user's authentication credentials.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Amazon Cognito uses the Amazon Simple Notification Service (SNS) to send silent push notifications to devices. A silent push notification is a push message that is received by your application on a user's device that will not be seen by the user.
Reference: http://aws.amazon.com/cognito/faqs/

QUESTION 289

Identify a correct statement about the expiration date of the "Letter of Authorization and Connecting Facility Assignment (LOA-CFA)," which lets you complete the Cross Connect step of setting up your AWS Direct Connect.

A. If the cross connect is not completed within 90 days, the authority granted by the LOA-CFA expires.
B. If the virtual interface is not created within 72 days, the LOA-CFA becomes outdated.
C. If the cross connect is not completed within a user-defined time, the authority granted by the LOA- CFA expires.
D. If the cross connect is not completed within the specified duration from the appropriate provider, the LOA-CFA expires.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
An AWS Direct Connect location provides access to AWS in the region it is associated with. You can establish connections with AWS Direct Connect locations in multiple regions, but a connection in one region does not provide connectivity to other regions. Note: If the cross connect is not completed within 90 days, the authority granted by the LOA-CFA expires.
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/Colocation.html

QUESTION 290
For AWS CloudFormation, which stack state refuses UpdateStack calls?

A. UPDATE_ROLLBACK_FAILED
B. UPDATE_ROLLBACK_COMPLETE
C. UPDATE_COMPLETE
D. CREATE_COMPLETE

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
When a stack is in the UPDATE_ROLLBACK_FAILED state, you can continue rolling it back to return it to a working state (to UPDATE_ROLLBACK_COMPLETE). You cannot update a stack that is in the UPDATE_ROLLBACK_FAILED state.
However, if you can continue to roll it back, you can return the stack to its original settings and try to update it again.
Reference: http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-updating-stacks-continueupdaterollback.html

QUESTION 291

A company runs an application on a fleet of Amazon EC2 instances. The application requires low latency and random access to 100 GB of data. The application must be able to access the data at up to 3.000 IOPS. A Development team has configured the EC2 launch template to provision a 100-GB Provisioned IOPS (PIOPS) Amazon EBS volume with 3 000 IOPS provisioned. A Solutions Architect is tasked with lowering costs without impacting performance and durability. Which action should be taken?

A. Create an Amazon EFS file system with the performance mode set to Max I/O. Configure the EC2 operating system to mount the EFS file system.
B. Create an Amazon EFS file system with the throughput mode set to Provisioned. Configure the EC2 operating system to mount the EFS file system.
C. Update the EC2 launch template to allocate a new 1-TB EBS General Purpose SSO (gp2) volume.
D. Update the EC2 launch template to exclude the PIOPS volume. Configure the application to use local instance storage.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 292
A company has a serverless application comprised of Amazon CloudFront, Amazon API Gateway, and AWS Lambda functions. The current deployment process of the application code is to create a new version number of the Lambda function and run an AWS CLI script to update. If the new function version has errors, another CLI script reverts by deploying the previous working version of the function. The company would like to decrease the time to deploy new versions of the application logic provided by the Lambda functions, and also reduce the time to detect and revert when errors are identified.
How can this be accomplished?

A. Create and deploy nested AWS CloudFormation stacks with the parent stack consisting of the AWS CloudFront distribution and API Gateway, and the child stack containing the Lambda function. For changes to Lambda, create an AWS CloudFormation change set and deploy; if errors are triggered, revert the AWS CloudFormation change set to the previous version.
B. Use AWS SAM and built-in AWS CodeDeploy to deploy the new Lambda version, gradually shift traffic to the new version, and use pre-traffic and post-traffic test functions to verify code. Rollback if Amazon CloudWatch alarms are triggered.
C. Refactor the AWS CLI scripts into a single script that deploys the new Lambda version. When deployment is completed, the script tests execute. If errors are detected, revert to the previous Lambda version.
D. Create and deploy an AWS CloudFormation stack that consists of a new API Gateway endpoint that references the new Lambda version. Change the CloudFront origin to the new API Gateway endpoint, monitor errors and if detected, change the AWS CloudFront origin to the previous API Gateway endpoint.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

Explanation:
Reference: https://aws.amazon.com/about-aws/whats-new/2017/11/aws-lambda-supports-traffic-shifting-and-phaseddeployments-with-aws-codedeploy/ https://docs.aws.amazon.com/serverless-applicationmodel/latest/developerguide/automating-updates-to-serverless-apps.html

QUESTION 293
With respect to AWS Lambda permissions model, at the time you create a Lambda function, you specify an IAM role that AWS Lambda can assume to execute your Lambda function on your behalf. This role is also referred to as the_____role.

A. configuration
B. execution
C. delegation
D. dependency

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Regardless of how your Lambda function is invoked, AWS Lambda always executes the function. At the time you create a Lambda function, you specify an IAM role that AWS Lambda can assume to execute your Lambda function on your behalf.
This role is also referred to as the execution role.
Reference: http://docs.aws.amazon.com/lambda/latest/dg/lambda-dg.pdf

QUESTION 294
By default, what is the maximum number of Cache Nodes you can run in Amazon ElastiCache?

A. 20
B. 50

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In Amazon ElastiCache, you can run a maximum of 20 Cache Nodes.

QUESTION 295
A company experienced a breach of highly confidential personal information due to permission issues on an Amazon S3 bucket. The Information Security team

has tightened the bucket policy to restrict access. Additionally, to be better prepared for future attacks, these requirements must be met:
Identify remote IP addresses that are accessing the bucket objects.
Receive alerts when the security policy on the bucket is changed. Remediate the policy changes automatically.
Which strategies should the Solutions Architect use?

A. Use Amazon CloudWatch Logs with CloudWatch filters to identify remote IP addresses. Use CloudWatch Events rules with AWS Lambda to automatically remediate S3 bucket policy changes. Use Amazon SES with CloudWatch Events rules for alerts.
B. Use Amazon Athena with S3 access logs to identify remote IP addresses. Use AWS Config rules with AWS Systems Manager Automation to automatically remediate S3 bucket policy changes. Use Amazon SNS with AWS Config rules for alerts.
C. Use S3 access logs with Amazon Elasticsearch Service and Kibana to identify remote IP addresses. Use an Amazon Inspector assessment template to automatically remediate S3 bucket policy changes. Use Amazon SNS for alerts.
D. Use Amazon Macie with an S3 bucket to identify access patterns and remote IP addresses. Use AWS Lambda with Macie to automatically remediate S3 bucket policy changes. Use Macie automatic alerting capabilities for alerts.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 296
A company uses an Amazon EMR cluster to process data once a day. The raw data comes from Amazon S3, and the resulting processed data is also stored in Amazon S3. The processing must complete within 4 hours; currently, it only takes 3 hours. However, the processing time is taking 5 to 10 minutes longer each week due to an increasing volume of raw data.
The team is also concerned about rising costs as the compute capacity increases. The EMR cluster is currently running on three m3.xlarge instances (one master and two core nodes).
Which of the following solutions will reduce costs related to the increasing compute needs?

A. Add additional task nodes, but have the team purchase an all-upfront convertible Reserved Instance for each additional node to offset the costs.
B. Add additional task nodes, but use instance fleets with the master node in on-Demand mode and a mix of On-Demand and Spot Instances for the core and task nodes. Purchase a scheduled Reserved Instance for the master node.
C. Add additional task nodes, but use instance fleets with the master node in Spot mode and a mix of On-Demand and Spot Instances for the core and task nodes. Purchase enough scheduled Reserved Instances to offset the cost of running any On- Demand instances.
D. Add additional task nodes, but use instance fleets with the master node in On-Demand mode and a mix of On-Demand and Spot Instances for the core and task nodes. Purchase a standard all-upfront Reserved Instance for the master node.

Correct Answer: B
Section: (none)

Explanation

Explanation/Reference:

QUESTION 297
A company is migrating its three-tier web application from on-premises to the AWS Cloud. The company has the following requirements for the migration process:
Ingest machine images from the on-premises environment.
Synchronize changes from the on-premises environment to the AWS environment until the production cutover.
Minimize downtime when executing the production cutover.
Migrate the virtual machines' root volumes and data volumes.
Which solution will satisfy these requirements with minimal operational overhead?

A. Use AWS Server Migration Service (SMS) to create and launch a replication job for each tier of the application. Launch instances from the AMIs created by AWS SMS. After initial testing, perform a final replication and create new instances from the updated AMIs.
B. Create an AWS CLI VM Import/Export script to migrate each virtual machine. Schedule the script to run incrementally to maintain changes in the application. Launch instances from the AMIs created by VM Import/Export. Once testing is done, rerun the script to do a final import and launch the instances from the AMIs.
C. Use AWS Server Migration Service (SMS) to upload the operating system volumes. Use the AWS CLI import-snapshot command for the data volumes. Launch instances from the AMIs created by AWS SMS and attach the data volumes to the instances. After initial testing, perform a final replication, launch new instances from the replicated AMIs, and attach the data volumes to the instances.
D. Use AWS Application Discovery Service and AWS Migration Hub to group the virtual machines as an application. Use the AWS CLI VM Import/Export script to import the virtual machines as AMIs. Schedule the script to run incrementally to maintain changes in the application. Launch instances from the AMIs. After initial testing, perform a final virtual machine import and launch new instances from the AMIs.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 298
In Amazon ElastiCache, the failure of a single cache node can have an impact on the availability of your application and the load on your back-end database while ElastiCache provisions a replacement for the failed cache node and it get repopulated.
Which of the following is a solution to reduce this potential availability impact?

A. Spread your memory and compute capacity over fewer number of cache nodes, each with smaller capacity.
B. Spread your memory and compute capacity over a larger number of cache nodes, each with smaller capacity.

C. Include fewer number of high capacity nodes.
D. Include a larger number of cache nodes, each with high capacity.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In Amazon ElastiCache, the number of cache nodes in the cluster is a key factor in the availability of your cluster running Memcached. The failure of a single cache node can have an impact on the availability of your application and the load on your back-end database while ElastiCache provisions a replacement for the failed cache node and it get repopulated.
You can reduce this potential availability impact by spreading your memory and compute capacity over a larger number of cache nodes, each with smaller capacity, rather than using a fewer number of high capacity nodes.
Reference: http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/CacheNode.Memcached.html

QUESTION 299
Identify an application that polls AWS Data Pipeline for tasks and then performs those tasks.

A. A task executor
B. A task deployer
C. A task runner
D. A task optimizer

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A task runner is an application that polls AWS Data Pipeline for tasks and then performs those tasks. You can either use Task Runner as provided by AWS Data Pipeline, or create a custom Task Runner application.
Task Runner is a default implementation of a task runner that is provided by AWS Data Pipeline. When Task Runner is installed and configured, it polls AWS Data Pipeline for tasks associated with pipelines that you have activated. When a task is assigned to Task Runner, it performs that task and reports its status back to AWS Data Pipeline. If your workflow requires non-default behavior, you'll need to implement that functionality in a custom task runner.
Reference: http://docs.aws.amazon.com/datapipeline/latest/DeveloperGuide/dp-how-remote-taskrunner-client.html

QUESTION 300
A retail company runs a business-critical web service on an Amazon Elastic Container Service (Amazon ECS) cluster that runs on Amazon EC2 instances. The web service receives POST requests from end users and writes data to a MySQL database that runs on a separate EC2 instance. The company needs to ensure

that data loss does not occur.
The current code deployment process includes manual updates of the ECS service. During a recent deployment, end users encountered intermittent 502 Bad Gateway errors in response to valid web requests.
The company wants to implement a reliable solution to prevent this issue from recurring. The company also wants to automate code deployments. The solution must be highly available and must optimize cost-effectiveness.
Which combination of steps will meet these requirements? (Choose three.)

A. Run the web service on an ECS cluster that has a Fargate launch type. Use AWS CodePipeline and AWS CodeDeploy to perform a blue/green deployment with validation testing to update the ECS service.
B. Migrate the MySQL database to run on an Amazon RDS for MySQL Multi-AZ DB instance that uses Provisioned IOPS SSD (io2) storage.
C. Configure an Amazon Simple Queue Service (Amazon SQS) queue as an event source to receive the POST requests from the web service. Configure an AWS Lambda function to poll the queue. Write the data to the database.
D. Run the web service on an ECS cluster that has a Fargate launch type. Use AWS CodePipeline and AWS CodeDeploy to perform a canary deployment to update the ECS service.
E. Configure an Amazon Simple Queue Service (Amazon SQS) queue. Install the SQS agent on the containers that run in the ECS cluster to poll the queue. Write the data to the database.
F. Migrate the MySQL database to run on an Amazon RDS for MySQL Multi-AZ DB instance that uses General Purpose SSD (gp3) storage.

Correct Answer: BCD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 301
Which EC2 functionality allows the user to place the Cluster Compute instances in clusters?

A. Cluster group
B. Cluster security group
C. GPU units
D. Cluster placement group

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

The Amazon EC2 cluster placement group functionality allows users to group cluster compute instances in clusters.
Reference: https://aws.amazon.com/ec2/faqs/

QUESTION 302
A company has an on-premises monitoring solution using a PostgreSQL database for persistence of events. The database is unable to scale due to heavy ingestion and it frequently runs out of storage.
The company wants to create a hybrid solution and has already set up a VPN connection between its network and AWS.
The solution should include the following attributes: Managed AWS services to minimize operational complexity.
A buffer that automatically scales to match the throughput of data and requires no ongoing administration.
A visualization tool to create dashboards to observe events in near-real time. Support for semi-structured JSON data and dynamic schemas.
Which combination of components will enable the company to create a monitoring solution that will satisfy these requirements? (Choose two.)

A. Use Amazon Kinesis Data Firehose to buffer events. Create an AWS Lambda function to process and transform events.
B. Create an Amazon Kinesis data stream to buffer events. Create an AWS Lambda function to process and transform events.
C. Configure an Amazon Aurora PostgreSQL DB cluster to receive events. Use Amazon QuickSight to read from the database and create near-real-time visualizations and dashboards.
D. Configure Amazon Elasticsearch Service (Amazon ES) to receive events. Use the Kibana endpoint deployed with Amazon ES to create near-real-time visualizations and dashboards.
E. Configure an Amazon Neptune DB instance to receive events. Use Amazon QuickSight to read from the database and create near-real-time visualizations and dashboards.

Correct Answer: BC
Section: (none)
Explanation

Explanation/Reference:

QUESTION 303
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 two.)

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: AD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 304
True or False: In Amazon ElastiCache replication groups of Redis, for performance tuning reasons, you can change the roles of the cache nodes within the replication group, with the primary and one of the replicas exchanging roles.

A. True, however, you get lower performance.
B. FALSE
C. TRUE
D. False, you must recreate the replication group to improve performance tuning.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In Amazon ElastiCache, a replication group is a collection of Redis Cache Clusters, with one primary read-write cluster and up to five secondary, read-only clusters, which are called read replicas. You can change the roles of the cache clusters within the replication group, with the primary cluster and one of the replicas exchanging roles. You might decide to do this for performance tuning reasons.
Reference: http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/Replication.Redis.Groups.html

QUESTION 305
A user is configuring MySQL RDS with PIOPS. What should be the minimum PIOPS that the user should provision?

A. 1000
B. 200
C. 2000
D. 500

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
If a user is trying to enable PIOPS with MySQL RDS, the minimum size of storage should be 100 GB and the minimum PIOPS should be 1000.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PIOPS.html

QUESTION 306
A company with several AWS accounts is using AWS Organizations and service control policies (SCPs). An Administrator created the following SCP and has attached it to an organizational unit (OU) that contains AWS account 1111-1111-1111:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowsAllActions",
            "Effect": "Allow",
            "Action": "*",
            "Resource": "*"
        },
        {
            "Sid": "DenyCloudTrail",
            "Effect": "Deny",
            "Action": "cloudtrail:*",
            "Resource": "*"
        }
    ]
}
```


Developers working in account 1111-1111-1111 complain that they cannot create Amazon S3 buckets. How should the Administrator address this problem?

A. Add s3:CreateBucket with "Allow" effect to the SCP.
B. Remove the account from the OU, and attach the SCP directly to account 1111-1111-1111.
C. Instruct the Developers to add Amazon S3 permissions to their IAM entities.
D. Remove the SCP from account 1111-1111-1111.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 307
A company currently uses Amazon EBS and Amazon RDS for storage purposes. The company intends to use a pilot light approach for disaster recovery in a different AWS Region. The company has an RTO of 6 hours and an RPO of 24 hours.
Which solution would achieve the requirements with MINIMAL cost?

A. Use AWS Lambda to create daily EBS and RDS snapshots, and copy them to the disaster recovery region. Use Amazon Route 53 with active-passive failover configuration. Use Amazon EC2 in an Auto Scaling group with the capacity set to 0 in the disaster recovery region.
B. Use AWS Lambda to create daily EBS and RDS snapshots, and copy them to the disaster recovery region. Use Amazon Route 53 with active-active failover configuration. Use Amazon EC2 in an Auto Scaling group configured in the same way as in the primary region.
C. Use Amazon ECS to handle long-running tasks to create daily EBS and RDS snapshots, and copy to the disaster recovery region. Use Amazon Route 53 with active-passive failover configuration. Use Amazon EC2 in an Auto Scaling group with the capacity set to 0 in the disaster recovery region.
D. Use EBS and RDS cross-region snapshot copy capability to create snapshots in the disaster recovery region. Use Amazon Route 53 with active-active failover configuration. Use Amazon EC2 in an Auto Scaling group with the capacity set to 0 in the disaster recovery region.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://amazonaws-china.com/about-aws/whats-new/2013/06/11/amazon-announces-faster-cross-region-ebs-snapshotcopy/

QUESTION 308
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 309
A mobile gaming company is expanding into the global market. The company's game servers run in the us-east-1 Region.
The game's client application uses UDP to communicate with the game servers and needs to be able to connect to a set of static IP addresses.
The company wants its game to be accessible on multiple continents. The company also wants the game to maintain its network performance and global availability.
Which solution meets these requirements?

A. Provision an Application Load Balancer (ALB) in front of the game servers. Create an Amazon CloudFront distribution that has no geographical restrictions. Set the ALB as the origin. Perform DNS lookups for the cloudfront.net domain name. Use the resulting IP addresses in the game's client application.
B. Provision game servers in each AWS Region. Provision an Application Load Balancer in front of the game servers. Create an Amazon Route 53 latency-based routing policy for the game's client application to use with DNS lookups.
C. Provision game servers in each AWS Region. Provision a Network Load Balancer (NLB) in front of the game servers. Create an accelerator in AWS Global Accelerator, and configure endpoint groups in each Region. Associate the NLBs with the corresponding Regional endpoint groups. Point the game client's application to the Global Accelerator endpoints.
D. Provision game servers in each AWS Region. Provision a Network Load Balancer (NLB) in front of the game servers. Create an Amazon CloudFront distribution that has no geographical restrictions. Set the NLB as the origin. Perform DNS lookups for the cloudfront.net domain name. Use the resulting IP addresses in the game's client application.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/global-accelerator/faqs/

QUESTION 310
To abide by industry regulations, a Solutions Architect must design a solution that will store a company's critical data in multiple public AWS Regions, including in the United States, where the company's headquarters is located. The Solutions Architect is required to provide access to the data stored in AWS to the company's global WAN network. The Security team mandates that no traffic accessing this data should traverse the public internet.
How should the Solutions Architect design a highly available solution that meets the requirements and is cost-effective?

A. Establish AWS Direct Connect connections from the company headquarters to all AWS Regions in use. Use the company WAN to send traffic over to the headquarters and then to the respective DX connection to access the data.
B. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connection. Use inter-region VPC peering to access the data in other AWS Regions.
C. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connection. Use an AWS transit VPC solution to access data in other AWS Regions.
D. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connection. Use Direct Connect Gateway to access data in other AWS Regions.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://aws.amazon.com/blogs/aws/new-aws-direct-connect-gateway-inter-region-vpc-access/

QUESTION 311
A solutions architect is troubleshooting an application that runs on Amazon EC2 instances. The EC2 instances run in an Auto Scaling group. The application needs to access user data in an Amazon DynamoDB table that has fixed provisioned capacity.
To match the increased workload, the solutions architect recently doubled the maximum size of the Auto Scaling group.
Now, when many instances launch at the same time, some application components are throttled when the components scan the DynamoDB table. The Auto Scaling group terminates the failing instances and starts new instances until all applications are running A solutions architect must implement a solution to mitigate the throttling issue in the MOST cost-effective manner Which solution will meet these requirements?

A. Double the provisioned read capacity of the DynamoDB table.
B. Duplicate the DynamoDB table. Configure the running copy of the application to select at random which table it access.
C. Set the DynamoDB table to on-demand mode.
D. Add DynamoDB Accelerator (DAX) to the table.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/premiumsupport/knowledge-center/on-demand-table-throttling-dynamodb/

QUESTION 312
A Solutions Architect has been asked to look at a company's Amazon Redshift cluster, which has quickly become an integral part of its technology and supports key business process. The Solutions Architect is to increase the reliability and availability of the cluster and provide options to ensure that if an issue arises, the cluster can either operate or be restored within four hours.
Which of the following solution options BEST addresses the business need in the most cost-effective manner?

A. Ensure that the Amazon Redshift cluster has been set up to make use of Auto Scaling groups with the nodes in the cluster spread across multiple Availability Zones.
B. Ensure that the Amazon Redshift cluster creation has been templated using AWS CloudFormation so it can easily be launched in another Availability Zone and data populated from the automated Redshift back-ups stored in Amazon S3.
C. Use Amazon Kinesis Data Firehose to collect the data ahead of ingestion into Amazon Redshift and create clusters using AWS CloudFormation in another region and stream the data to both clusters.
D. Create two identical Amazon Redshift clusters in different regions (one as the primary, one as the secondary). Use Amazon S3 cross-region replication from the primary to secondary region, which triggers an AWS Lambda function to populate the cluster in the secondary region.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Q: What happens to my data warehouse cluster availability and data durability if my data warehouse cluster's Availability Zone (AZ) has an outage?
If your Amazon Redshift data warehouse cluster's Availability Zone becomes unavailable, you will not be able to use your cluster until power and network access to the AZ are restored. Your data warehouse cluster's data is preserved so you can start using your Amazon Redshift data warehouse as soon as the AZ becomes available again. In addition, you can also choose to restore any existing snapshots to a new AZ in the same Region. Amazon Redshift will restore your most frequently accessed data first so you can resume queries as quickly as possible.
Reference: https://aws.amazon.com/redshift/faqs/?nc1=h_ls

QUESTION 313
A company built an application based on AWS Lambda deployed in an AWS CloudFormation stack. The last production release of the web application introduced an issue that resulted in an outage lasting several minutes. A solutions architect must adjust the deployment process to support a canary release.
Which solution will meet these requirements?

A. Create an alias for every new deployed version of the Lambda function. Use the AWS CLI update-alias command with the routing-config parameter to distribute the load.
B. Deploy the application into a new CloudFormation stack. Use an Amazon Route 53 weighted routing policy to distribute the load.
C. Create a version for every new deployed Lambda function. Use the AWS CLI update-function-configuration command with the routing-config parameter to distribute the load.
D. Configure AWS CodeDeploy and use CodeDeployDefault.OneAtATime in the Deployment configuration to distribute the load.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 314
A Solutions Architect must establish a patching plan for a large mixed fleet of Windows and Linux servers. The patching plan must be implemented securely, be audit-ready, and comply with the company's business requirements.
Which option will meet these requirements with MINIMAL effort?

A. Install and use an OS-native patching service to manage the update frequency and release approval for all instances. Use AWS Config to verify the OS state on each instance and report on any patch compliance issues.
B. Use AWS Systems Manager on all instances to manage patching. Test patches outside of production and then deploy during a maintenance window with the appropriate approval.
C. Use AWS OpsWorks for Chef Automate to run a set of scripts that will iterate through all instances of a given type. Issue the appropriate OS command to get and install updates on each instance, including any required restarts during the maintenance window.
D. Migrate all applications to AWS OpsWorks and use OpsWorks automatic patching support to keep the OS up-to-date following the initial installation. Use AWS Config to provide audit and compliance reporting.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Only Systems Manager can patch both OS effectively on AWS and on premise.

QUESTION 315
What combination of steps could a Solutions Architect take to protect a web workload running on Amazon EC2 from DDoS and application layer attacks? (Choose two.)

A. Put the EC2 instances behind a Network Load Balancer and configure AWS WAF on it.
B. Migrate the DNS to Amazon Route 53 and use AWS Shield.
C. Put the EC2 instances in an Auto Scaling group and configure AWS WAF on it.
D. Create and use an Amazon CloudFront distribution and configure AWS WAF on it.
E. Create and use an internet gateway in the VPC and use AWS Shield.

Correct Answer: DE
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/answers/networking/aws-ddos-attack-mitigation/

QUESTION 316
A financial company is using a high-performance compute cluster running on Amazon EC2 instances to perform market simulations. A DNS record must be created in an Amazon Route 53 private hosted zone when instances start. The DNS record must be removed after instances are terminated.
Currently the company uses a combination of Amazon CloudWatch Events and AWS Lambda to create the DNS record. The solution worked well in testing with small clusters, but in production with clusters containing thousands of instances the company sees the following error in the Lambda logs:
HTTP 400 error (Bad request).
The response header also includes a status code element with a value of “Throttling” and a status message element with a value of “Rate exceeded”.
Which combination of steps should the Solutions Architect take to resolve these issues? (Choose three.)

A. Configure an Amazon SOS FIFO queue and configure a CloudWatch Events rule to use this queue as a target. Remove the Lambda target from the CloudWatch Events rule.
B. Configure an Amazon Kinesis data stream and configure a CloudWatch Events rule to use this queue as a target. Remove the Lambda target from the CloudWatch Events rule.
C. Update the CloudWatch Events rule to trigger on Amazon EC2 “Instance Launch Successful” and “Instance Terminate Successful” events for the Auto Scaling group used by the cluster.
D. Configure a Lambda function to retrieve messages from an Amazon SQS queue. Modify the Lambda function to retrieve a maximum of 10 messages then batch the messages by Amazon Route 53 API call type and submit. Delete the messages from the SQS queue after successful API calls.
E. Configure an Amazon SQS standard queue and configure the existing CloudWatch Events rule to use this queue as a target. Remove the Lambda target from the CloudWatch Events rule.
F. Configure a Lambda function to read data from the Amazon Kinesis data stream and configure the batch window to 5 minutes. Modify the function to make a single API call to Amazon Route 53 with all records read from the kinesis data stream.

Correct Answer: BEF
Section: (none)
Explanation

Explanation/Reference:

QUESTION 317
Regarding Identity and Access Management (IAM), Which type of special account belonging to your application allows your code to access Google services

programmatically?

A. Service account
B. Simple Key
C. OAuth
D. Code account

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A service account is a special Google account that can be used by applications to access Google services programmatically.
This account belongs to your application or a virtual machine (VM), instead of to an individual end user. Your application uses the service account to call the Google API of a service, so that the users aren't directly involved.
A service account can have zero or more pairs of service account keys, which are used to authenticate to Google. A service account key is a public/private key pair generated by Google. Google retains the public key, while the user is given the private key.
Reference:
https://cloud.google.com/iam/docs/service-accounts

QUESTION 318
A company is migrating its data center from on premises to the AWS Cloud. The migration will take several months to complete. The company will use Amazon Route 53 for private DNS zones.
During the migration, the company must keep its AWS services pointed at the VPC's Route 53 Resolver for DNS. The company also must maintain the ability to resolve addresses from its on-premises DNS server. A solutions architect must set up DNS so that Amazon EC2 instances can use native Route 53 endpoints to resolve on-premises DNS queries.
Which configuration will meet these requirements?

A. Configure the VPC DHCP options set to point to on-premises DNS server IP addresses. Ensure that security groups for EC2 instances allow outbound access to port 53 on those DNS server IP addresses.
B. Launch an EC2 instance that has DNS BIND installed and configured. Ensure that the security groups that are attached to the EC2 instance can access the on-premises DNS server IP address on port 53. Configure BIND to forward DNS queries to on-premises DNS server IP addresses. Configure each migrated EC2 instance's DNS settings to point to the BIND server IP address.
C. Create a new outbound endpoint in Route 53, and attach the endpoint to the VPEnsure that the security groups that are attached to the endpoint can access the on-premises DNS server IP address on port 53. Create a new Route 53 Resolver rule that routes on-premises designated traffic to the on-premises DNS server.
D. Create a new private DNS zone in Route 53 with the same domain name as the on-premises domain. Create a single wildcard record with the on-premises DNS server IP address as the record's address.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver.html

QUESTION 319
A solutions architect needs to deploy an application on a fleet of Amazon EC2 Instances. The EC2 instances run in private subnets in an Auto Scaling group. The application is expected to generate logs at a rate of 100 MB each second on each of the EC2 instances.
The logs must be stored in an Amazon S3 bucket so that an Amazon EMR cluster can consume them for further processing.
The logs must be quickly accessible for the first 90 days and should be retrievable within 48 hours thereafter.
What is the MOST cost-effective solution that meets these requirements?

A. Set up an S3 copy job to write logs from each EC2 instance to the S3 bucket with S3 Standard storage. Use a NAT instance within the private subnets to connect to Amazon S3. Create S3 Lifecycle policies to move logs that are older than 90 days to S3 Glacier.
B. Set up an S3 sync job to copy logs from each EC2 instance to the S3 bucket with S3 Standard storage. Use a gateway VPC endpoint for Amazon S3 to connect to Amazon S3. Create S3 Lifecycle policies to move logs that are older than 90 days to S3 Glacier Deep Archive.
C. Set up an S3 batch operation to copy logs from each EC2 instance to the S3 bucket with S3 Standard storage. Use a NAT gateway with the private subnets to connect to Amazon S3. Create S3 Lifecycle policies to move logs that are older than 90 days to S3 Glacier Deep Archive.
D. Set up an S3 sync job to copy logs from each EC2 instance to the S3 bucket with S3 Standard storage. Use a gateway VPC endpoint for Amazon S3 to connect to Amazon S3. Create S3 Lifecycle policies to move logs that are older than 90 days to S3 Glacier.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 320
A company has developed a new billing application that will be released in two weeks. Developers are testing the application running on 10 EC2 instances managed by an Auto Scaling group in subnet 172.31.0.0/24 within VPC A with CIDR block 172.31.0.0/16. The Developers noticed connection timeout errors in the application logs while connecting to an Oracle database running on an Amazon EC2 instance in the same region within VPC B with CIDR block 172.50.0.0/16. The IP of the database instance is hard-coded in the application instances.
Which recommendations should a Solutions Architect present to the Developers to solve the problem in a secure way with minimal maintenance and overhead?

A. Disable the SrcDestCheck attribute for all instances running the application and Oracle Database. Change the default route of VPC A to point ENI of the Oracle Database that has an IP address assigned within the range of 172.50.0.0/16

B. Create and attach internet gateways for both VPCs. Configure default routes to the internet gateways for both VPCs. Assign an Elastic IP for each Amazon EC2 instance in VPC A

C. Create a VPC peering connection between the two VPCs and add a route to the routing table of VPC A that points to the IP address range of 172.50.0.0/16

D. Create an additional Amazon EC2 instance for each VPC as a customer gateway; create one virtual private gateway (VGW) for each VPC, configure an end-to-end VPC, and advertise the routes for 172.50.0.0/16

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 321
An enterprise company is building an infrastructure services platform for its users. The company has the following requirements:
Provide least privilege access to users when launching AWS infrastructure so users cannot provision unapproved services.
Use a central account to manage the creation of infrastructure services.
Provide the ability to distribute infrastructure services to multiple accounts in AWS Organizations. Provide the ability to enforce tags on any infrastructure that is started by users.
Which combination of actions using AWS services will meet these requirements? (Choose three.)

A. Develop infrastructure services using AWS Cloud Formation templates. Add the templates to a central Amazon S3 bucket and add the-IAM roles or users that require access to the S3 bucket policy.

B. Develop infrastructure services using AWS Cloud Formation templates. Upload each template as an AWS Service Catalog product to portfolios created in a central AWS account. Share these portfolios with the Organizations structure created for the company.

C. Allow user IAM roles to have AWSCloudFormationFullAccess and AmazonS3ReadOnlyAccess permissions. Add an Organizations SCP at the AWS account root user level to deny all services except AWS CloudFormation and Amazon S3.

D. Allow user IAM roles to have ServiceCatalogEndUserAccess permissions only. Use an automation script to import the central portfolios to local AWS accounts, copy the TagOption assign users access and apply launch constraints.

E. Use the AWS Service Catalog TagOption Library to maintain a list of tags required by the company. Apply the TagOption to AWS Service Catalog products or portfolios.

F. Use the AWS CloudFormation Resource Tags property to enforce the application of tags to any CloudFormation templates that will be created for users.

Correct Answer: ABE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 322
A user is hosting a public website on AWS. The user wants to have the database and the app server on the AWS VPC. The user wants to setup a database that can connect to the Internet for any patch upgrade but cannot receive any request from the internet. How can the user set this up?

A. Setup DB in a private subnet with the security group allowing only outbound traffic.
B. Setup DB in a public subnet with the security group allowing only inbound data.
C. Setup DB in a local data center and use a private gateway to connect the application with DB.
D. Setup DB in a private subnet which is connected to the internet via NAT for outbound.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. AWS provides two features that the user can use to increase security in VPC: security groups and network ACLs. When the user wants to setup both the DB and App on VPC, the user should make one public and one private subnet. The DB should be hosted in a private subnet and instances in that subnet cannot reach the internet. The user can allow an instance in his VPC to initiate outbound connections to the internet but prevent unsolicited inbound connections from the internet by using a Network Address Translation (NAT) instance.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

QUESTION 323
A company's factory and automation applications are running in a single VPC. More than 20 applications run on a combination of Amazon EC2, Amazon Elastic Container Service (Amazon ECS), and Amazon RDS.
The company has software engineers spread across three teams. One of the three teams owns each application, and each time is responsible for the cost and performance of all of its applications. Team resources have tags that represent their application and team. The teams use IAM access for daily activities.
The company needs to determine which costs on the monthly AWS bill are attributable to each application or team. The company also must be able to create reports to compare costs from the last 12 months and to help forecast costs for the next 12 months. A solutions architect must recommend an AWS Billing and Cost Management solution that provides these cost reports. Which combination of actions will meet these requirements? (Choose three.)

A. Activate the user-define cost allocation tags that represent the application and the team.
B. Activate the AWS generated cost allocation tags that represent the application and the team.
C. Create a cost category for each application in Billing and Cost Management.
D. Activate IAM access to Billing and Cost Management.
E. Create a cost budget.
F. Enable Cost Explorer.

Correct Answer: BCF
Section: (none)
Explanation

Explanation/Reference:

QUESTION 324
A company provides AWS solutions to its users with AWS CloudFormation templates. Users launch the templates in their accounts to have different solutions provisioned for them. The users want to improve the deployment strategy for solutions while retaining the ability to do the following:
Add their own features to a solution for their specific deployments.
Run unit tests on their changes.
Turn features on and off for their deployments.
Automatically update with code changes.
Run security scanning tools for their deployments.
Which strategies should the Solutions Architect use to meet the requirements?

A. Allow users to download solution code as Docker images. Use AWS CodeBuild and AWS CodePipeline for the CI/CD pipeline. Use Docker images for different solution features and the AWS CLI to turn features on and off. Use AWS CodeDeploy to run unit tests and security scans, and for deploying and updating a solution with changes.
B. Allow users to download solution code artifacts. Use AWS CodeCommit and AWS CodePipeline for the CI/CD pipeline. Use AWS Amplify plugins for different solution features and user prompts to turn features on and off. Use AWS Lambda to run unit tests and security scans, and AWS CodeBuild for deploying and updating a solution with changes.
C. Allow users to download solution code artifacts in their Amazon S3 buckets. Use Amazon S3 and AWS CodePipeline for the CI/CD pipelines. Use CloudFormation StackSets for different solution features and to turn features on and off. Use AWS Lambda to run unit tests and security scans, and CloudFormation for deploying and updating a solution with changes.
D. Allow users to download solution code artifacts. Use AWS CodeCommit and AWS CodePipeline for the CI/CD pipeline. Use the AWS Cloud Development Kit constructs for different solution features, and use the manifest file to turn features on and off. Use AWS CodeBuild to run unit tests and security scans, and for deploying and updating a solution with changes.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://www.slideshare.net/AmazonWebServices/cicd-for-containers-a-way-forward-for-your-devops-pipeline

QUESTION 325
A company has multiple AWS accounts as part of an organization created with AWS Organizations. Each account has a VPC in the us-east-2 Region and is

used for either production or development workloads. Amazon EC2 instances across production accounts need to communicate with each other, and EC2 instances across development accounts need to communicate with each other, but production and development instances should not be able to communicate with each other.
To facilitate connectivity, the company created a common network account. The company used AWS Transit Gateway to create a transit gateway in the us-east-2 Region in the network account and shared the transit gateway with the entire organization by using AWS Resource Access Manager. Network administrators then attached VPCs in each account to the transit gateway, after which the EC2 instances were able to communicate across accounts. However, production and development accounts were also able to communicate with one another.
Which set of steps should a solutions architect take to ensure production traffic and development traffic are completely isolated?

A. Modify the security groups assigned to development EC2 instances to block traffic from production EC2 instances. Modify the security groups assigned to production EC2 instances to block traffic from development EC2 instances.
B. Create a tag on each VPC attachment with a value of either production or development, according to the type of account being attached. Using the Network Manager feature of AWS Transit Gateway, create policies that restrict traffic between VPCs based on the value of this tag.
C. Create separate route tables for production and development traffic. Delete each account's association and route propagation to the default AWS Transit Gateway route table. Attach development VPCs to the development AWS Transit Gateway route table and production VPCs to the production route table, and enable automatic route propagation on each attachment.
D. Create a tag on each VPC attachment with a value of either production or development, according to the type of account being attached. Modify the AWS Transit Gateway routing table to route production tagged attachments to one another and development tagged attachments to one another.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/vpc/latest/tgw/tgw-transit-gateways.html

QUESTION 326
A software company hosts an application on AWS with resources in multiple AWS accounts and Regions. The application runs on a group of Amazon EC2 instances in an application VPC located in the useast-1 Region with an IPv4 CIDR block of 10.10.0.0/16. In a different AWS account, a shared services VPC is located in the us-east-2 Region with an IPv4 CIDR block of 10.10.10.0/24. When a cloud engineer uses AWS CloudFormation to attempt to peer the application VPC with the shared services VPC, an error message indicates a peering failure. Which factors could cause this error? (Choose two.)

A. The IPv4 CIDR ranges of the two VPCs overlap
B. The VPCs are not in the same Region
C. One or both accounts do not have access to an Internet gateway
D. One of the VPCs was not shared through AWS Resource Access Manager
E. The IAM role in the peer accepter account does not have the correct permissions

Correct Answer: AE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 327
A company has an application that sells tickets online and experiences bursts of demand every 7 days. The application has a stateless presentation layer running on Amazon EC2, an Oracle database to store unstructured data catalog information, and a backend API layer. The front-end layer uses an Elastic Load Balancer to distribute the load across nine On-Demand instances over three Availability Zones (AZs). The Oracle database is running on a single EC2 instance. The company is experiencing performance issues when running more than two concurrent campaigns. A solutions architect must design a solution that meets the following requirements:
Address scalability issues.
Increase the level of concurrency.
Eliminate licensing costs. Improve reliability.
Which set of steps should the solutions architect take?

A. Create an Auto Scaling group for the front end with a combination of On-Demand and Spot Instances to reduce costs. Convert the Oracle database into a single Amazon RDS reserved DB instance.
B. Create an Auto Scaling group for the front end with a combination of On-Demand and Spot Instances to reduce costs. Create two additional copies of the database instance, then distribute the databases in separate AZs.
C. Create an Auto Scaling group for the front end with a combination of On-Demand and Spot Instances to reduce costs. Convert the tables in the Oracle database into Amazon DynamoDB tables.
D. Convert the On-Demand Instances into Spot instances to reduce costs for the front end. Convert the tables in the Oracle database into Amazon DynamoDB tables.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 328
A solutions architect needs to define a reference architecture for a solution for three-tier applications with web, application, and NoSQL data layers. The reference architecture must meet the following requirements:
High availability within an AWS Region
Able to fail over in 1 minute to another AWS Region for disaster recovery Provide the most efficient solution while minimizing the impact on the user experience
Which combination of steps will meet these requirements? (Choose three.)

A. Use an Amazon Route 53 weighted routing policy set to 100/0 across the two selected Regions. Set Time to Live (TTL) to 1 hour.
B. Use an Amazon Route 53 failover routing policy for failover from the primary Region to the disaster recovery Region. Set Time to Live (TTL) to 30 seconds.
C. Use a global table within Amazon DynamoDB so data can be accessed in the two selected Regions.
D. Back up data from an Amazon DynamoDB table in the primary Region every 60 minutes and then write the data to Amazon S3. Use S3 cross-Region replication to copy the data from the primary Region to the disaster recovery Region.
Have a script import the data into DynamoDB in a disaster recovery scenario.
E. Implement a hot standby model using Auto Scaling groups for the web and application layers across multiple Availability Zones in the Regions. Use zonal Reserved Instances for the minimum number of servers and On-Demand Instances for any additional resources.
F. Use Auto Scaling groups for the web and application layers across multiple Availability Zones in the Regions. Use Spot Instances for the required resources.

Correct Answer: ACE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 329
A company runs a Windows Server host in a public subnet that is configured to allow a team of administrators to connect over RDP to troubleshoot issues with hosts in a private subnet. The host must be available at all times outside of a scheduled maintenance window, and needs to receive the latest operating system updates within 3 days of release.
What should be done to manage the host with the LEAST amount of administrative effort?

A. Run the host in a single-instance AWS Elastic Beanstalk environment. Configure the environment with a custom AMI to use a hardened machine image from AWS Marketplace. Apply system updates with AWS Systems Manager Patch Manager.
B. Run the host on AWS WorkSpaces. Use Amazon WorkSpaces Application Manager (WAM) to harden the host. Configure Windows automatic updates to occur every 3 days.
C. Run the host in an Auto Scaling group with a minimum and maximum instance count of 1. Use a hardened machine image from AWS Marketplace. Apply system updates with AWS Systems Manager Patch Manager.
D. Run the host in AWS OpsWorks Stacks. Use a Chief recipe to harden the AMI during instance launch. Use an AWS Lambda scheduled event to run the Upgrade Operating System stack command to apply system updates.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Reference: https://docs.aws.amazon.com/workspaces/latest/adminguide/workspace-maintenance.html

QUESTION 330
A European online newspaper service hosts its public-facing WordPress site in a collocated data center in London. The current WordPress infrastructure consists of a load balancer, two web servers, and one MySQL database server. A solutions architect is tasked with designing a solution with the following requirements: Improve the website's performance Make the web tier scalable and stateless Improve the database server performance for read-heavy loads Reduce latency for users across Europe and the US Design the new architecture with a goal of 99.9% availability Which solution meets these requirements while optimizing operational efficiency?

A. Use an Application Load Balancer (ALB) in front of an Auto Scaling group of WordPress Amazon EC2 instances in one AWS Region and three Availability Zones. Configure an Amazon ElastiCache cluster in front of a Multi-AZ Amazon Aurora MySQL DB cluster. Move the WordPress shared files to Amazon EFS. Configure Amazon CloudFront with the ALB as the origin, and select a price class that includes the US and Europe.
B. Use an Application Load Balancer (ALB) in front of an Auto Scaling group of WordPress Amazon EC2 instances in two AWS Regions and two Availability Zones in each Region. Configure an Amazon ElastiCache cluster in front of a global Amazon Aurora MySQL database. Move the WordPress shared files to Amazon EFS. Configure Amazon CloudFront with the ALB as the origin, and select a price class that includes the US and Europe. Configure EFS cross-Region replication.
C. Use an Application Load Balancer (ALB) in front of an Auto Scaling group of WordPress Amazon EC2 instances in one AWS Region and three Availability Zones. Configure an Amazon DocumentDB table in front of a Multi-AZ Amazon Aurora MySQL DB cluster. Move the WordPress shared files to Amazon EFS. Configure Amazon CloudFront with the ALB as the origin, and select a price class that includes all global locations.
D. Use an Application Load Balancer (ALB) in front of an Auto Scaling group of WordPress Amazon EC2 instances in two AWS Regions and three Availability Zones in each Region. Configure an Amazon ElastiCache cluster in front of a global Amazon Aurora MySQL database. Move the WordPress shared files to Amazon FSx with cross-Region synchronization. Configure Amazon CloudFront with the ALB as the origin and a price class that includes the US and Europe.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 331
A large mobile gaming company has successfully migrated all of its on-premises infrastructure to the AWS Cloud. A solutions architect is reviewing the environment to ensure that it was built according to the design and that it is running in alignment with the Well-Architected Framework.
While reviewing previous monthly costs in Cost Explorer, the solutions architect notices that the creation and subsequent termination of several large instance types account for a high proportion of the costs. The solutions architect finds out that the company's developers are launching new Amazon EC2 instances as part of their testing and that the developers are not using the appropriate instance types.
The solutions architect must implement a control mechanism to limit the instance types that only the developers can launch.
Which solution will meet these requirements?

A. Create a desired-instance-type managed rule in AWS Config. Configure the rule with the instance types that are allowed.

Attach the rule to an event to run each time a new EC2 instance is launched.

B. In the EC2 console, create a launch template that specifies the instance types that are allowed. Assign the launch template to the developers' IAM accounts.

C. Create a new IAM policy. Specify the instance types that are allowed. Attach the policy to an IAM group that contains the IAM accounts for the developers

D. Use EC2 Image Builder to create an image pipeline for the developers and assist them in the creation of a golden image.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/config/latest/developerguide/evaluate-config_develop-rules_getting-started.html

QUESTION 332
A company is using an Amazon CloudFront distribution to distribute both static and dynamic content from a web application running behind an Application Load Balancer. The web application requires user authorization and session tracking for dynamic content. The CloudFront distribution has a single cache behavior configured to forward the Authorization, Host, and User-Agent HTTP whitelist headers and a session cookie to the origin. All other cache behavior settings are set to their default value.
A valid ACM certificate is applied to the CloudFront distribution with a matching CNAME in the distribution settings. The ACM certificate is also applied to the HTTPS listener for the Application Load Balancer. The CloudFront origin protocol policy is set to HTTPS only. Analysis of the cache statistics report shows that the miss rate for this distribution is very high.
What can the Solutions Architect do to improve the cache hit rate for this distribution without causing the SSL/TLS handshake between CloudFront and the Application Load Balancer to fail?

A. Create two cache behaviors for static and dynamic content. Remove the User-Agent and Host HTTP headers from thewhitelist headers section on both of the cache behaviors. Remove the session cookie from the whitelist cookies section andthe Authorization HTTP header from the whitelist headers section for cache behavior configured for static content.

B. Remove the User-Agent and Authorization HTTP headers from the whitelist headers section of the cache behavior. Thenupdate the cache behavior to use presigned cookies for authorization.

C. Remove the Host HTTP header from the whitelist headers section and remove the session cookie from the whitelistcookies section for the default cache behavior. Enable automatic object compression and use Lambda@Edge viewerrequest events for user authorization.

D. Create two cache behaviors for static and dynamic content. Remove the User-Agent HTTP header from the whitelistheaders section on both of the cache behaviors. Remove the session cookie from the whitelist cookies section and theAuthorization HTTP header from the whitelist headers section for cache behavior configured for static content.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 333
An AWS partner company is building a service in AWS Organizations using its organization named org1. This service requires the partner company to have access to AWS resources in a customer account, which is in a separate organization named org2. The company must establish least privilege security access using an API or command line tool to the customer account.
What is the MOST secure way to allow org1 to access resources in org2?

A. The customer should provide the partner company with their AWS account access keys to log in and perform the required tasks.
B. The customer should create an IAM user and assign the required permissions to the IAM user. The customer should then provide the credentials to the partner company to log in and perform the required tasks.
C. The customer should create an IAM role and assign the required permissions to the IAM role. The partner company should then use the IAM role's Amazon Resource Name (ARN) when requesting access to perform the required tasks.
D. The customer should create an IAM role and assign the required permissions to the IAM role. The partner company should then use the IAM role's Amazon Resource Name (ARN), including the external ID in the IAM role's trust policy, when requesting access to perform the required tasks.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 334
A Solutions Architect is working with a company that is extremely sensitive to its IT costs and wishes to implement controls that will result in a predictable AWS spend each month.
Which combination of steps can help the company control and monitor its monthly AWS usage to achieve a cost that is as close as possible to the target amount? (Choose three.)

A. Implement an IAM policy that requires users to specify a 'workload' tag for cost allocation when launching Amazon EC2 instances.
B. Contact AWS Support and ask that they apply limits to the account so that users are not able to launch more than a certain number of instance types.
C. Purchase all upfront Reserved Instances that cover 100% of the account's expected Amazon EC2 usage.
D. Place conditions in the users' IAM policies that limit the number of instances they are able to launch.
E. Define 'workload' as a cost allocation tag in the AWS Billing and Cost Management console.
F. Set up AWS Budgets to alert and notify when a given workload is expected to exceed a defined cost.

Correct Answer: AEF
Section: (none)

Explanation

Explanation/Reference:

QUESTION 335
A car rental company has built a serverless REST API to provide data to its mobile app. The app consists of an Amazon API Gateway API with a Regional endpoint, AWS Lambda functions, and an Amazon Aurora MySQL Serverless DB cluster. The company recently opened the API to mobile apps of partners. A significant increase in the number of requests resulted, causing sporadic database memory errors. Analysis of the API traffic indicates that clients are making multiple HTTP GET requests for the same queries in a short period of time. Traffic is concentrated during business hours, with spikes around holidays and other events.
The company needs to improve its ability to support the additional usage while minimizing the increase in costs associated with the solution.
Which strategy meets these requirements?

A. Convert the API Gateway Regional endpoint to an edge-optimized endpoint. Enable caching in the production stage.
B. Implement an Amazon ElastiCache for Redis cache to store the results of the database calls. Modify the Lambda functions to use the cache.
C. Modify the Aurora Serverless DB cluster configuration to increase the maximum amount of available memory.
D. Enable throttling in the API Gateway production stage. Set the rate and burst values to limit the incoming calls.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/getting-started/projects/build-serverless-web-app-lambda-apigateway-s3-dynamodbcognito/module-4/

QUESTION 336
Out of the striping options available for the EBS volumes, which one has the following disadvantage:
'Doubles the amount of I/O required from the instance to EBS compared to RAID 0, because you're mirroring all writes to a pair of volumes, limiting how much you can stripe.'?

A. Raid 1
B. Raid 0
C. RAID 1+0 (RAID 10)
D. Raid 2

Correct Answer: C
Section: (none)

Explanation

Explanation/Reference:
Explanation:
RAID 1+0 (RAID 10) doubles the amount of I/O required from the instance to EBS compared to RAID 0, because you're mirroring all writes to a pair of volumes, limiting how much you can stripe.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/raid-config.html

QUESTION 337
Which of the following cannot be done using AWS Data Pipeline?

A. Create complex data processing workloads that are fault tolerant, repeatable, and highly available.
B. Regularly access your data where it's stored, transform and process it at scale, and efficiently transfer the results to another AWS service.
C. Generate reports over data that has been stored.
D. Move data between different AWS compute and storage services as well as on premise data sources at specified intervals.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
AWS Data Pipeline is a web service that helps you reliably process and move data between different AWS compute and storage services as well as on premise data sources at specified intervals. With AWS Data Pipeline, you can regularly access your data where it's stored, transform and process it at scale, and efficiently transfer the results to another AWS.
AWS Data Pipeline helps you easily create complex data processing workloads that are fault tolerant, repeatable, and highly available. AWS Data Pipeline also allows you to move and process data that was previously locked up in on premise data silos.
Reference: http://aws.amazon.com/datapipeline/

QUESTION 338
The following AWS Identity and Access Management (IAM) customer managed policy has been attached to an IAM user:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
        "Effect": "Allow",
        "Action": "s3:*",
        "Resource": [
          "arn:aws:s3:::prod-data",
          "arn:aws:s3:::prod-data/*"
        ]
    },
    {
        "Effect":"Deny",
        "NotAction": "s3:*",
        "NotResource": [
          "arn:aws:s3:::prod-data",
          "arn:aws:s3:::prod-data/*"
        ]
    }
  ]
}
```

Which statement describes the access that this policy provides to the user?

A. The policy grants access to all Amazon S3 actions, including all actions in the prod-data S3 bucket
B. This policy denies access to all Amazon S3 actions, excluding all actions in the prod-data S3 bucket
C. This policy denies access to the Amazon S3 bucket and objects not having prod-data in the bucket name
D. This policy grants access to all Amazon S3 actions in the prod-data S3 bucket, but explicitly denies access to all other AWS services

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 339
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 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: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_enable-console-saml.html

QUESTION 340
A company is currently using AWS CodeCommit for its source control and AWS CodePipeline for continuous integration.
The pipeline has a build stage for building the artifacts, which is then staged in an Amazon S3 bucket.
The company has identified various improvement opportunities in the existing process, and a Solutions Architect has been given the following requirements:
Create a new pipeline to support feature development
Support feature development without impacting production applications Incorporate continuous testing with unit tests Isolate development and production artifacts Support the capability to merge tested code into production code.
How should the Solutions Architect achieve these requirements?

A. Trigger a separate pipeline from CodeCommit feature branches. Use AWS CodeBuild for running unit tests. Use CodeBuild to stage the artifacts within an S3 bucket in a separate testing account.
B. Trigger a separate pipeline from CodeCommit feature branches. Use AWS Lambda for running unit tests. Use AWS CodeDeploy to stage the artifacts within an S3 bucket in a separate testing account.
C. Trigger a separate pipeline from CodeCommit tags. Use Jenkins for running unit tests. Create a stage in the pipeline with S3 as the target for staging the artifacts with an S3 bucket in a separate testing account.
D. Create a separate CodeCommit repository for feature development and use it to trigger the pipeline. Use AWS Lambda for running unit tests. Use AWS CodeBuild to stage the artifacts within different S3 buckets in the same production account.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

Explanation:
Reference:
https://docs.aws.amazon.com/codebuild/latest/userguide/how-to-create-pipeline.html

QUESTION 341
An organization is purchasing licensed software. The software license can be registered only to a specific MAC Address.
The organization is going to host the software in the AWS environment.
How can the organization fulfil the license requirement as the MAC address changes every time an instance is started/stopped/terminated?

A. It is not possible to have a fixed MAC address with AWS.
B. The organization should use VPC with the private subnet and configure the MAC address with that subnet.
C. The organization should use VPC with an elastic network interface which will have a fixed MAC Address.
D. The organization should use VPC since VPC allows to configure the MAC address for each EC2 instance.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. An Elastic Network Interface (ENI) is a virtual network interface that the user can attach to an instance in a VPC. An ENI can include attributes such as: a primary private IP address, one or more secondary private IP addresses, one elastic IP address per private IP address, one public IP address, one or more security groups, a MAC address, a source/destination check flag, and a description. The user can create a network interface, attach it to an instance, detach it from an instance, and attach it to another instance. The attributes of a network interface follow the network interface as it is attached or detached from an instance and reattached to another instance. Thus, the user can maintain a fixed MAC using the network interface.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html

QUESTION 342
A company is running an application distributed over several Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer. The security team requires that all application access attempts be made available for analysis.
Information about the client IP address, connection type, and user agent must be included.
Which solution will meet these requirements?

A. Enable EC2 detailed monitoring, and include network logs. Send all logs through Amazon Kinesis Data Firehose to an Amazon Elasticsearch Service (Amazon ES) cluster that the security team uses for analysis.
B. Enable VPC Flow Logs for all EC2 instance network interfaces. Publish VPC Flow Logs to an Amazon S3 bucket. Have the security team use Amazon Athena to query and analyze the logs.
C. Enable access logs for the Application Load Balancer, and publish the logs to an Amazon S3 bucket. Have the security team use Amazon Athena to query and analyze the logs.

D. Enable Traffic Mirroring and specify all EC2 instance network interfaces as the source. Send all traffic information through Amazon Kinesis Data Firehose to an Amazon Elasticsearch Service (Amazon ES) cluster that the security team uses for analysis.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 343
A solutions architect is designing a disaster recovery strategy for a three-tier application. The application has an RTO of 30 minutes and an RPO of 5 minutes for the data tier. The application and web tiers are stateless and leverage a fleet of Amazon EC2 instances. The data tier consists of a 50 TB Amazon Aurora database.
Which combination of steps satisfies the RTO and RPO requirements while optimizing costs? (Choose two.)

A. Create daily snapshots of the EC2 instances and replicate the snapshots to another Region.
B. Deploy a hot standby of the application to another Region.
C. Create snapshots of the Aurora database every 5 minutes.
D. Create a cross-Region Aurora Replica of the database.
E. Create an AWS Backup job to replicate data to another Region.

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 344
Once the user has set ElastiCache for an application and it is up and running, which services, does Amazon not provide for the user:

A. The ability for client programs to automatically identify all of the nodes in a cache cluster, and to initiate and maintain connections to all of these nodes
B. Automating common administrative tasks such as failure detection and recovery, and software patching.
C. Providing default Time to Live (TTL) in the AWS Elasticache Redis Implementation for different type of data.
D. Providing detailed monitoring metrics associated with your Cache Nodes, enabling you to diagnose and react to issues very quickly

Correct Answer: C
Section: (none)

Explanation

Explanation/Reference:
Explanation:
Amazon provides failure detection and recovery, and software patching and monitoring tools which is called CloudWatch. In addition it provides also Auto Discovery to automatically identify and initialize all nodes of cache cluster for Amazon ElastiCache.
Reference: http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/WhatIs.html

QUESTION 345
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 all rule.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Organizations usually implement proxy solutions to provide URL and web content filtering, IDS/IPS, data loss prevention, monitoring, and advanced threat protection. Reference: https://d0.awsstatic.com/awsanswers/ Controlling_VPC_Egress_Traffic.pdf

QUESTION 346
A company wants to migrate its website from an on-premises data center onto AWS. At the same time, it wants to migrate the website to a containerized microservice-based architecture to improve the availability and cost efficiency. The company's security policy states that privileges and network permissions must be configured according to best practice, using least privilege.
A Solutions Architect must create a containerized architecture that meets the security requirements and has deployed the application to an Amazon ECS cluster.
What steps are required after the deployment to meet the requirements? (Choose two.)

A. Create tasks using the bridge network mode.
B. Create tasks using the awsvpc network mode.
C. Apply security groups to Amazon EC2 instances, and use IAM roles for EC2 instances to access other resources.

D. Apply security groups to the tasks, and pass IAM credentials into the container at launch time to access other resources.
E. Apply security groups to the tasks, and use IAM roles for tasks to access other resources.

Correct Answer: BE
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/about-aws/whats-new/2017/11/amazon-ecs-introduces-awsvpc-networking-mode-forcontainers-to-support-full-networking-capabilities/ https://amazonaws-china.com/blogs/compute/introducing-cloud-nativenetworking-for-ecs-containers/ https://docs.aws.amazon.com/AmazonECS/latest/developerguide/task-iam-roles.html

QUESTION 347
A company has five physical data centers in specific locations around the world. Each data center has hundreds of physical servers with a mix of Windows and Linux-based applications and database services. Each data center also has an AWS Direct Connect connection of 10 Gbps to AWS with a company-approved VPN solution to ensure that data transfer is secure.
The company needs to shut down the existing data centers as quickly as possible and migrate the servers and applications to AWS.
Which solution meets these requirements?

A. Install the AWS Server Migration Service (AWS SMS) connector onto each physical machine. Use the AWS Management Console to select the servers from the server catalog, and start the replication. Once the replication is complete, launch the Amazon EC2 instances created by the service.
B. Install the AWS DataSync agent onto each physical machine. Use the AWS Management Console to configure the destination to be an AMI, and start the replication. Once the replication is complete, launch the Amazon EC2 instances created by the service.
C. Install the CloudEndure Migration agent onto each physical machine. Create a migration blueprint, and start the replication. Once the replication is complete, launch the Amazon EC2 instances in cutover mode.
D. Install the AWS Application Discovery Service agent onto each physical machine. Use the AWS Migration Hub import option to start the replication. Once the replication is complete, launch the Amazon EC2 instances created by the service.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 348
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:
Explanation:
Reference: https://aws.amazon.com/directconnect/faqs/

QUESTION 349
A company currently has data hosted in an IBM Db2 database. A web application calls an API that runs stored procedures on the database to retrieve user information data that is read-only. This data is historical in nature and changes on a daily basis. When a user logs in to the application, this data needs to be retrieved within 3 seconds. Each time a user logs in, the stored procedures run. Users log in several times a day to check stock prices.
Running this database has become cost-prohibitive due to Db2 CPU licensing. Performance goals are not being met.
Timeouts from Db2 are common due to long-running queries.
Which approach should a solutions architect take to migrate this solution to AWS?

A. Rehost the Db2 database in Amazon Fargate. Migrate all the data. Enable caching in Fargate. Refactor the API to use the Fargate Db2 database. Implement Amazon API Gateway and enable API caching.
B. Use AWS DMS to migrate data to Amazon DynamoDB using a continuous replication task. Refactor the API to use the DynamoDB data. Implement the refactored API in Amazon API Gateway and enable API caching.
C. Create a local cache on the mainframe to store query outputs. Use SFTP to sync to Amazon S3 on a daily basis. Refactor the API to use Amazon EFS. Implement Amazon API Gateway and enable API caching.
D. Extract data daily and copy the data to AWS Snowball for storage on Amazon S3. Sync daily. Refactor the API to use the S3 data. Implement Amazon API Gateway and enable API caching.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 350
A company is using an existing orchestration tool to manage thousands of Amazon EC2 instances. A recent penetration test found a vulnerability in the company's software stack. This vulnerability has prompted the company to perform a full evaluation of its current production environment. The analysis determined that the following vulnerabilities exist within the environment:
Operating systems with outdated libraries and known vulnerabilities are being used in production.
Relational databases hosted and managed by the company are running unsupported versions with known vulnerabilities.
Data stored in databases is not encrypted.
The solutions architect intends to use AWS Config to continuously audit and assess the compliance of the company's AWS resource configurations with the company's policies and guidelines.
What additional steps will enable the company to secure its environments and track resources while adhering to best practices?

A. Use AWS Application Discovery Service to evaluate all running EC2 instances Use the AWS CLI to modify each instance, and use EC2 user data to install the AWS Systems Manager Agent during boot. Schedule patching to run as a Systems Manager Maintenance Windows task. Migrate all relational databases to Amazon RDS and enable AWS KMS encryption.
B. Create an AWS CloudFormation template for the EC2 instances. Use EC2 user data in the CloudFormation template to install the AWS Systems Manager Agent, and enable AWS KMS encryption on all Amazon EBS volumes. Have CloudFormation replace all running instances. Use Systems Manager Patch Manager to establish a patch baseline and deploy a Systems Manager Maintenance Windows task to execute AWS-RunPatchBaseline using the patch baseline.
C. Install the AWS Systems Manager Agent on all existing instances using the company's current orchestration tool. Use the Systems Manager Run Command to execute a list of commands to upgrade software on each instance using operating system-specific tools. Enable AWS KMS encryption on all Amazon EBS volumes.
D. Install the AWS Systems Manager Agent on all existing instances using the company's current orchestration tool. Migrate all relational databases to Amazon RDS and enable AWS KMS encryption. Use Systems Manager Patch Manager to establish a patch baseline and deploy a Systems Manager Maintenance Windows task to execute AWS-RunPatchBaseline using the patch baseline.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 351
A financial services company sells its software-as-a-service (SaaS) platform for application compliance to large global banks.
The SaaS platform runs on AWS and uses multiple AWS accounts that are managed in an organization in AWS Organizations. The SaaS platform uses many AWS resources globally.
For regulatory compliance, all API calls to AWS resources must be audited, tracked for changes, and stored in a durable and secure data store.

Which solution will meet these requirements with the LEAST operational overhead?

A. Create a new AWS CloudTrail trail. Use an existing Amazon S3 bucket in the organization's management account to store the logs. Deploy the trail to all AWS Regions. Enable MFA delete and encryption on the S3 bucket.
B. Create a new AWS CloudTrail trail in each member account of the organization. Create new Amazon S3 buckets to store the logs. Deploy the trail to all AWS Regions. Enable MFA delete and encryption on the S3 buckets.
C. Create a new AWS CloudTrail trail in the organization's management account. Create a new Amazon S3 bucket with versioning turned on to store the logs. Deploy the trail for all accounts in the organization. Enable MFA delete and encryption on the S3 bucket.
D. Create a new AWS CloudTrail trail in the organization's management account. Create a new Amazon S3 bucket to store the logs. Configure Amazon Simple Notification Service (Amazon SNS) to send log-file delivery notifications to an external management system that will track the logs. Enable MFA delete and encryption on the S3 bucket.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/awscloudtrail/latest/userguide/creating-an-organizational-trail-in-the-console.html

QUESTION 352
The Solutions Architect manages a serverless application that consists of multiple API gateways, AWS Lambda functions, Amazon S3 buckets, and Amazon DynamoDB tables. Customers say that a few application components slow while loading dynamic images, and some are timing out with the "504 Gateway Timeout" error. While troubleshooting the scenario, the Solutions Architect confirms that DynamoDB monitoring metrics are at acceptable levels.
Which of the following steps would be optimal for debugging these application issues? (Choose two.)

A. Parse HTTP logs in Amazon API Gateway for HTTP errors to determine the root cause of the errors.
B. Parse Amazon CloudWatch Logs to determine processing times for requested images at specified intervals.
C. Parse VPC Flow Logs to determine if there is packet loss between the Lambda function and S3.
D. Parse AWS X-Ray traces and analyze HTTP methods to determine the root cause of the HTTP errors.
E. Parse S3 access logs to determine if objects being accessed are from specific IP addresses to narrow the scope to geographic latency issues.

Correct Answer: BD
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Firstly "A 504 Gateway Timeout Error means your web server didn't receive a timely response from another server upstream when it attempted to load one of

your web pages. Put simply, your web servers aren't communicating with each other fast enough". This specific issue is addressed in the AWS article "Tracing, Logging and Monitoring an API Gateway API".
Reference:
https://docs.amazonaws.cn/en_us/apigateway/latest/developerguide/monitoring_overview.html

QUESTION 353
An education company is running a web application used by college students around the world. The application runs in an Amazon Elastic Container Service (Amazon ECS) cluster in an Auto Scaling group behind an Application Load Balancer (ALB). A system administrator detects a weekly spike in the number of failed login attempts, which overwhelm the application's authentication service. All the failed login attempts originate from about 500 different IP addresses that change each week. A solutions architect must prevent the failed login attempts from overwhelming the authentication service.
Which solution meets these requirements with the MOST operational efficiency?

A. Use AWS Firewall Manager to create a security group and security group policy to deny access from the IP addresses
B. Create an AWS WAF web ACL with a rate-based rule, and set the rule action to Block. Connect the web ACL to the ALB
C. Use AWS Firewall Manager to create a security group and security group policy to allow access only to specific CIDR ranges
D. Create an AWS WAF web ACL with an IP set match rule, and set the rule action to Block. Connect the web ACL to the ALB

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/waf/latest/developerguide/security-group-policies.html

QUESTION 354
A company is migrating a legacy application from an on-premises data center to AWS. The application consists of a single application server and a Microsoft SQL Server database server. Each server is deployed on a VMware VM that consumes 500 TB of data across multiple attached volumes.
The company has established a 10 Gbps AWS Direct Connect connection from the closest AWS Region to its on-premises data center. The Direct Connect connection is not currently in use by other services.
Which combination of steps should a solutions architect take to migrate the application with the LEAST amount of downtime?
(Choose two.)

A. Use an AWS Server Migration Service (AWS SMS) replication job to migrate the database server VM to AWS.
B. Use VM Import/Export to import the application server VM.
C. Export the VM images to an AWS Snowball Edge Storage Optimized device.
D. Use an AWS Server Migration Service (AWS SMS) replication job to migrate the application server VM to AWS.
E. Use an AWS Database Migration Service (AWS DMS) replication instance to migrate the database to an Amazon RDS DB instance.

Correct Answer: BE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 355
A company has an Amazon EC2 deployment that has the following architecture:
An application tier that contains 8 m4.xlarge instances
A Classic Load Balancer
Amazon S3 as a persistent data store
After one of the EC2 instances fails, users report very slow processing of their requests. A Solutions Architect must recommend design changes to maximize system reliability. The solution must minimize costs.
What should the Solutions Architect recommend?

A. Migrate the existing EC2 instances to a serverless deployment using AWS Lambda functions
B. Change the Classic Load Balancer to an Application Load Balancer
C. Replace the application tier with m4.large instances in an Auto Scaling group
D. Replace the application tier with 4 m4.2xlarge instances

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
By default, connection draining is enabled for Application Load Balancers but must be enabled for Classic Load Balancers.
When Connection Draining is enabled and configured, the process of deregistering an instance from an Elastic Load Balancer gains an additional step. For the duration of the configured timeout, the load balancer will allow existing, in-flight requests made to an instance to complete, but it will not send any new requests to the instance. During this time, the API will report the status of the instance as InService, along with a message stating that “Instance deregistration currently in progress.” Once the timeout is reached, any remaining connections will be forcibly closed.
Reference: https://docs.aws.amazon.com/autoscaling/ec2/userguide/attach-load-balancer-asg.html
https://aws.amazon.com/blogs/aws/elb-connection-draining-remove-instances-from-service-with-care/

QUESTION 356
A medical company is building a data lake on Amazon S3. The data must be encrypted in transit and at rest. The data must remain protected even if S3 bucket is inadvertently made public.
Which combination of steps will meet these requirements? (Choose three.)

A. Ensure that each S3 bucket has a bucket policy that includes a Deny statement if the aws:SecureTransport condition is not present.
B. Create a CMK in AWS Key Management Service (AWS KMS). Turn on server-side encryption (SSE) on the S3 buckets, select SSE-KMS for the encryption type, and use the CMK as the key.
C. Ensure that each S3 bucket has a bucket policy that includes a Deny statement for PutObject actions if the request does not include an “s3:x-amz-server-side-encryption”:“aws:kms” condition.
D. Turn on server-side encryption (SSE) on the S3 buckets and select SSE-S3 for the encryption type.
E. Ensure that each S3 bucket has a bucket policy that includes a Deny statement for PutObject actions if the request does not include an “s3:x-amz-server-side-encryption”:“AES256” condition.
F. Turn on AWS Config. Use the s3-bucket-public-read-prohibited, s3-bucket-public-write-prohibited, and s3-bucket-sslrequests- only AWS Config managed rules to monitor the S3 buckets.

Correct Answer: ABC
Section: (none)
Explanation

Explanation/Reference:
Explanation:
To determine HTTP or HTTPS requests in a bucket policy, use a condition that checks for the key "aws:SecureTransport".
When this key is true, then request is sent through HTTPS. To comply with the s3bucket-ssl-requests-only rule, create abucket policy that explicitly denies access when the request meets the condition "aws:SecureTransport": "false". This policyexplicitly denies access to HTTP requests.
When you create an object, you can specify the use of server-side encryption with AWS Key Management Service (AWS KMS) keys to encrypt your data. This is true when you are either uploading a new object or copying an existing object. This encryption is known as SSE-KMS.
Enforce object encryption, create an S3 bucket policy that denies any S3 Put request that does not include the x-amz-serverside- encryption header.
Reference: https://aws.amazon.com/premiumsupport/knowledge-center/s3-bucket-policy-for-config-rule/
https://docs.aws.amazon.com/AmazonS3/latest/userguide/specifying-kms-encryption.html
https://aws.amazon.com/blogs/security/how-to-prevent-uploads-of-unencrypted-objects-to-amazon-s3/

QUESTION 357
A video streaming company recently launched a mobile app for video sharing. The app uploads various files to an Amazon S3 bucket in the us-east-1 Region. The files range in size from 1 GB to 10 GB.
Users who access the app from Australia have experienced uploads that take long periods of time. Sometimes the files fail to completely upload for these users.
A solutions architect must improve the app’s performance for these uploads.
Which solutions will meet these requirements? (Choose two.)

A. Enable S3 Transfer Acceleration on the S3 bucket. Configure the app to use the Transfer Acceleration endpoint for uploads.
B. Configure an S3 bucket in each Region to receive the uploads. Use S3 Cross-Region Replication to copy the files to the distribution S3 bucket.
C. Set up Amazon Route 53 with latency-based routing to route the uploads to the nearest S3 bucket Region.
D. Configure the app to break the video files into chunks. Use a multipart upload to transfer files to Amazon S3.
E. Modify the app to add random prefixes to the files before uploading.

Correct Answer: AC
Section: (none)
Explanation

Explanation/Reference:

QUESTION 358
If you have a running instance using an Amazon EBS boot partition, you can call the _____ API to release the compute resources but preserve the data on the boot partition.

A. Stop Instances
B. Terminate Instances
C. AMI Instance
D. Ping Instance

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
If you have a running instance using an Amazon EBS boot partition, you can also call the Stop Instances API to release the compute resources but preserve the data on the boot partition.
Reference: https://aws.amazon.com/ec2/faqs/#How_quickly_will_systems_be_running

QUESTION 359
An elastic network interface (ENI) is a virtual network interface that you can attach to an instance in a VPC. An ENI can include one public IP address, which can be auto-assigned to the elastic network interface for eth0 when you launch an instance, but only when you_____.

A. create an elastic network interface for eth1
B. include a MAC address
C. use an existing network interface
D. create an elastic network interface for eth0

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
An elastic network interface (ENI) is defined as a virtual network interface that you can attach to an instance in a VPC and can include one public IP address, which can be auto-assigned to the elastic network interface for eth0 when you launch an instance, but only when you create an elastic network interface for eth0 instead of using an existing network interface.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html

QUESTION 360
A user has enabled detailed CloudWatch monitoring with the AWS Simple Notification Service. Which of the below mentioned statements helps the user understand detailed monitoring better?

A. SNS cannot provide data every minute
B. SNS will send data every minute after configuration
C. There is no need to enable since SNS provides data every minute
D. AWS CloudWatch does not support monitoring for SNS

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
CloudWatch is used to monitor AWS as well as the custom services. It provides either basic or detailed monitoring for the supported AWS products. In basic monitoring, a service sends data points to CloudWatch every five minutes, while in detailed monitoring a service sends data points to CloudWatch every minute. The AWS SNS service sends data every 5 minutes. Thus, it supports only the basic monitoring. The user cannot enable detailed monitoring with SNS.
Reference: http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/supported_services.html

QUESTION 361
Which of the following commands accepts binary data as parameters?

A. --user-data
B. –cipher text-key
C. --aws-customer-key
D. --describe-instances-user

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
For commands that take binary data as a parameter, specify that the data is binary content by using the fileb:// prefix.
Commands that accept binary data include: aws ec2 run-instances --user-data parameter. aws s3api put-object --ssecustomer- key parameter. aws kms decrypt --ciphertext-blob parameter.
Reference: http://docs.aws.amazon.com/cli/latest/userguide/aws-cli.pdf

QUESTION 362
A company runs a video processing platform. Files are uploaded by users who connect to a web server, which stores them on an Amazon EFS share. This web server is running on a single Amazon EC2 instance. A different group of instances, running in an Auto Scaling group, scans the EFS share directory structure for new files to process and generates new videos (thumbnails, different resolution, compression, etc.) according to the instructions file, which is uploaded along with the video files. A different application running on a group of instances managed by an Auto Scaling group processes the video files and then deletes them from the EFS share. The results are stored in an S3 bucket. Links to the processed video files are emailed to the customer.
The company has recently discovered that as they add more instances to the Auto Scaling Group, many files are processed twice, so image processing speed is not improved. The maximum size of these video files is 2GB.
What should the Solutions Architect do to improve reliability and reduce the redundant processing of video files?

A. Modify the web application to upload the video files directly to Amazon S3. Use Amazon CloudWatch Events to trigger an AWS Lambda function every time a file is uploaded, and have this Lambda function put a message into an Amazon SQS queue. Modify the video processing application to read from SQS queue for new files and use the queue depth metric to scale instances in the video processing Auto Scaling group.
B. Set up a cron job on the web server instance to synchronize the contents of the EFS share into Amazon S3. Trigger an AWS Lambda function every time a file is uploaded to process the video file and store the results in Amazon S3. Using Amazon CloudWatch Events, trigger an Amazon SES job to send an email to the customer containing the link to the processed file.
C. Rewrite the web application to run directly from Amazon S3 and use Amazon API Gateway to upload the video files to an S3 bucket. Use an S3 trigger to run an AWS Lambda function each time a file is uploaded to process and store new video files in a different bucket. Using CloudWatch Events, trigger an SES job to send an email to the customer containing the link to the processed file.
D. Rewrite the web application to run from Amazon S3 and upload the video files to an S3 bucket. Each time a new file is uploaded, trigger an AWS Lambda function to put a message in an SQS queue containing the link and the instructions.
Modify the video processing application to read from the SQS queue and the S3 bucket. Use the queue depth metric to adjust the size of the Auto Scaling group for video processing instances.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/FIFO-queues.html

QUESTION 363
A company is running a commercial Apache Hadoop cluster on Amazon EC2. This cluster is being used daily to query large files on Amazon S3. The data on Amazon S3 has been curated and does not require any additional transformations steps.
The company is using a commercial business intelligence (BI) tool on Amazon EC2 to run queries against the Hadoop cluster and visualize the data.
The company wants to reduce or eliminate the overhead costs associated with managing the Hadoop cluster and the BI tool.
The company would like to move to a more cost-effective solution with minimal effort. The visualization is simple and requires performing some basic aggregation steps only.
Which option will meet the company's requirements?

A. Launch a transient Amazon EMR cluster daily and develop an Apache Hive script to analyze the files on Amazon S3.
Shut down the Amazon EMR cluster when the job is complete. Then use Amazon QuickSight to connect to Amazon EMR and perform the visualization.
B. Develop a stored procedure invoked from a MySQL database running on Amazon EC2 to analyze the files in Amazon S3.
Then use a fast in-memory BI tool running on Amazon EC2 to visualize the data.
C. Develop a script that uses Amazon Athena to query and analyze the files on Amazon S3. Then use Amazon QuickSight to connect to Athena and perform the visualization.
D. Use a commercial extract, transform, load (ETL) tool that runs on Amazon EC2 to prepare the data for processing. Then switch to a faster and cheaper BI tool that runs on Amazon EC2 to visualize the data from Amazon S3.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/quicksight/latest/user/create-a-data-set-athena.html
https://aws.amazon.com/athena/

QUESTION 364
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 data.
Which of these solutions would you recommend?

A. 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.
B. 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.
C. 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.

D. 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 365
How does in-memory caching improve the performance of applications in ElastiCache?

A. It improves application performance by deleting the requests that do not contain frequently accessed data.
B. It improves application performance by implementing good database indexing strategies.
C. It improves application performance by using a part of instance RAM for caching important data.
D. It improves application performance by storing critical pieces of data in memory for low-latency access.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In Amazon ElastiCache, in-memory caching improves application performance by storing critical pieces of data in memory for low-latency access. Cached information may include the results of I/Ointensive database queries or the results of computationally intensive calculations.
Reference: http://aws.amazon.com/elasticache/faqs/#g4

QUESTION 366
You want to mount an Amazon EFS file system on an Amazon EC2 instance using DNS names. Which of the following generic form of a mount target's DNS name must you use to mount the file system?

A. availability-zone.file-system-id.efs.aws-region.amazonaws.com
B. efs-system-id.availability-zone.file-aws-region.amazonaws.com
C. $file-system-id.$availability-zone.$efs.aws-region.$amazonaws.com
D. #aws-region.#availability-zone.#file-system-id.#efs.#amazonaws.com

Correct Answer: A
Section: (none)

Explanation

Explanation/Reference:
Explanation:
An Amazon EFS file system can be mounted on an Amazon EC2 instance using DNS names. This can be done with either a DNS name for the file system or a DNS name for the mount target. To construct the mount target's DNS name, use the following generic form: availability-zone.file-system-id.efs.aws-region.amazonaws.com
Reference: http://docs.aws.amazon.com/efs/latest/ug/mounting-fs.html#mounting-fs-install-nfsclient

QUESTION 367
You want to use AWS CodeDeploy to deploy an application to Amazon EC2 instances running within an Amazon Virtual Private Cloud (VPC).
What criterion must be met for this to be possible?

A. The AWS CodeDeploy agent installed on the Amazon EC2 instances must be able to access only the public AWS CodeDeploy endpoint.
B. The AWS CodeDeploy agent installed on the Amazon EC2 instances must be able to access only the public Amazon S3 service endpoint.
C. The AWS CodeDeploy agent installed on the Amazon EC2 instances must be able to access the public AWS CodeDeploy and Amazon S3 service endpoints.
D. It is not currently possible to use AWS CodeDeploy to deploy an application to Amazon EC2 instances running within an Amazon Virtual Private Cloud (VPC.)

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
You can use AWS CodeDeploy to deploy an application to Amazon EC2 instances running within an Amazon Virtual Private Cloud (VPC).
However, the AWS CodeDeploy agent installed on the Amazon EC2 instances must be able to access the public AWS CodeDeploy and Amazon S3 service endpoints. Reference: http://aws.amazon.com/codedeploy/faqs/

QUESTION 368
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 369
A company is running multiple workloads in the AWS Cloud. The company has separate units for software development. The company uses AWS Organizations and federation with SAML to give permissions to developers to manage resources in their AWS accounts. The development units each deploy their production workloads into a common production account.
Recently, an incident occurred in the production account in which members of a development unit terminated an EC2 instance that belonged to a different development unit. A solutions architect must create a solution that prevents a similar incident from happening in the future. The solution also must allow developers the possibility to manage the instances used for their workloads.
Which strategy will meet these requirements?

A. Create separate OUs in AWS Organizations for each development unit. Assign the created OUs to the company AWS accounts. Create separate SCPs with a deny action and a StringNotEquals condition for the DevelopmentUnit resource tag that matches the development unit name. Assign the SCP to the corresponding OU.
B. Pass an attribute for DevelopmentUnit as an AWS Security Token Service (AWS STS) session tag during SAML federation. Update the IAM policy for the developers' assumed IAM role with a deny action and a StringNotEquals condition for the DevelopmentUnit resource tag and aws:PrincipalTag/ DevelopmentUnit.
C. Pass an attribute for DevelopmentUnit as an AWS Security Token Service (AWS STS) session tag during SAML federation. Create an SCP with an allow action and a StringEquals condition for the DevelopmentUnit resource tag and aws:PrincipalTag/DevelopmentUnit. Assign the SCP to the root OU.
D. Create separate IAM policies for each development unit. For every IAM policy, add an allow action and a StringEquals condition for the DevelopmentUnit resource tag and the development unit name. During SAML federation, use AWS Security Token Service (AWS STS) to assign the IAM policy and match the development unit name to the assumed IAM role.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 370
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 6 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 upto- date standby. Compare this to a Single-AZ deployment: in case of a Single-AZ database failure, a user-initiated point-intime- 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 371
In the context of AWS IAM, identify a true statement about user passwords (login profiles).

A. They must contain Unicode characters.
B. They can contain any Basic Latin (ASCII) characters.
C. They must begin and end with a forward slash (/).
D. They cannot contain Basic Latin (ASCII) characters.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The user passwords (login profiles) of IAM users can contain any Basic Latin (ASCII)characters.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/LimitationsOnEntities.html

QUESTION 372
A company manages hundreds of AWS accounts centrally in an organization in AWS Organizations. The company recently started to allow product teams to create and manage their own S3 access points in their accounts. The S3 access points can be accessed only within VPCs, not on the Internet.
What is the MOST operationally efficient way to enforce this requirement?

A. Set the S3 access point resource policy to deny the s3:CreateAccessPoint action unless the s3:AccessPointNetworkOrigin condition key evaluates to VPC.
B. Create an SCP at the root level in the organization to deny the s3:CreateAccessPoint action unless the s3:AccessPointNetworkOrigin condition key evaluates to VPC.
C. Use AWS CloudFormation StackSets to create a new IAM policy in each AWS account that allows the s3:CreateAccessPoint action only if the s3:AccessPointNetworkOrigin condition key evaluates to VPC.

D. Set the S3 bucket policy to deny the s3:CreateAccessPoint action unless the s3:AccessPointNetworkOrigin condition key evaluates to VPC.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/blogs/storage/managing-amazon-s3-access-with-vpc-endpoints-and-s3-access-points/

QUESTION 373
A software company has deployed an application that consumes a REST API by using Amazon API Gateway, AWS Lambda functions, and an Amazon DynamoDB table. The application is showing an increase in the number of errors during PUT requests. Most of the PUT calls come from a small number of clients that are authenticated with specific API keys.
A solutions architect has identified that a large number of the PUT requests originate from one client. The API is noncritical, and clients can tolerate retries of unsuccessful calls. However, the errors are displayed to customers and are causing damage to the API's reputation.
What should the solutions architect recommend to improve the customer experience?

A. Implement retry logic with exponential backoff and irregular variation in the client application. Ensure that the errors are caught and handled with descriptive error messages.
B. Implement API throttling through a usage plan at the API Gateway level. Ensure that the client application handles code 429 replies without error.
C. Turn on API caching to enhance responsiveness for the production stage. Run 10-minute load tests. Verify that the cache capacity is appropriate for the workload.
D. Implement reserved concurrency at the Lambda function level to provide the resources that are needed during sudden increases in traffic.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
API Gateway recommends that you run a 10-minute load test to verify that your cache capacity is appropriate for your workload.
Reference: https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-caching.html

QUESTION 374
In which step of "start using AWS Direct Connect" steps is the virtual interface you created tagged with a customer-provided tag that complies with the Ethernet 802.1Q standard?

A. Download Router Configuration.
B. Complete the Cross Connect.

C. Configure Redundant Connections with AWS Direct Connect.
D. Create a Virtual Interface.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In the list of using Direct Connect steps, the create a Virtual Interface step is to provision your virtual interfaces. Each virtual interface must be tagged with a customer-provided tag that complies with the Ethernet 802.1Q standard. This tag is required for any traffic traversing the AWS Direct Connect connection.
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/getstarted.html#createvirtualinterface

QUESTION 375
A company has a data lake in Amazon S3 that needs to be accessed by hundreds of applications across many AWS accounts. The company's information security policy states that the S3 bucket must not be accessed over the public internet and that each application should have the minimum permissions necessary to function.
To meet these requirements, a solutions architect plans to use an S3 access point that is restricted to specific VPCs for each application.
Which combination of steps should the solutions architect take to implement this solution? (Choose two.)

A. Create an S3 access point for each application in the AWS account that owns the S3 bucket. Configure each access point to be accessible only from the application's VPC. Update the bucket policy to require access from an access point
B. Create an interface endpoint for Amazon S3 in each application's VPC. Configure the endpoint policy to allow access to an S3 access point. Create a VPC gateway attachment for the S3 endpoint
C. Create a gateway endpoint for Amazon S3 in each application's VPConfigure the endpoint policy to allow access to an S3 access point. Specify the route table that is used to access the access point.
D. Create an S3 access point for each application in each AWS account and attach the access points to the S3 bucket.
Configure each access point to be accessible only from the application's VPC. Update the bucket policy to require access from an access point.
E. Create a gateway endpoint for Amazon S3 in the data lake's VPC. Attach an endpoint policy to allow access to the S3 bucket. Specify the route table that is used to access the bucket

Correct Answer: AC
Section: (none)
Explanation

Explanation/Reference:

QUESTION 376
What is the network performance offered by the c4.8xlarge instance in Amazon EC2?

A. Very High but variable
B. 20 Gigabit
C. 5 Gigabit
D. 10 Gigabit

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Networking performance offered by the c4.8xlarge instance is 10 Gigabit.
Reference: http://aws.amazon.com/ec2/instance-types/

QUESTION 377
What is the default maximum number of VPCs allowed per region?

A. 5
B. 10
C. 100
D. 15

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The maximum number of VPCs allowed per region is 5.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Appendix_Limits.html

QUESTION 378
A company is using AWS for production and development workloads. Each business unit has its own AWS account for production, and a separate AWS account to develop and deploy its applications. The Information Security department has introduced new security policies that limit access for terminating certain Amazon EC2 instances in all accounts to a small group of individuals from the Security team.
How can the Solutions Architect meet these requirements?

A. Create a new IAM policy that allows access to those EC2 instances only for the Security team. Apply this policy to the AWS Organizations master account.

B. Create a new tag-based IAM policy that allows access to these EC2 instances only for the Security team. Tag the instances appropriately, and apply this policy in each account.
C. Create an organizational unit under AWS Organizations. Move all the accounts into this organizational unit and use SCP to apply a whitelist policy to allow access to these EC2 instances for the Security team only.
D. Set up SAML federation for all accounts in AWS. Configure SAML so that it checks for the service API call before authenticating the user. Block SAML from authenticating API calls if anyone other than the Security team accesses these instances.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://aws.amazon.com/blogs/security/how-to-use-service-control-policies-to-set-permission-guardrails-across-accounts-inyour-aws-organization/ https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_examplescps.html

QUESTION 379
A Solutions Architect is designing a multi-account structure that has 10 existing accounts. The design must meet the following requirements:
Consolidate all accounts into one organization.
Allow full access to the Amazon EC2 service from the master account and the secondary accounts. Minimize the effort required to add additional secondary accounts.
Which combination of steps should be included in the solution? (Choose two.)

A. Create an organization from the master account. Send invitations to the secondary accounts from the master account. Accept the invitations and create an OU.
B. Create an organization from the master account. Send a join request to the master account from each secondary account. Accept the requests and create an OU.
C. Create a VPC peering connection between the master account and the secondary accounts. Accept the request for the VPC peering connection.
D. Create a service control policy (SCP) that enables full EC2 access, and attach the policy to the OU.
E. Create a full EC2 access policy and map the policy to a role in each account. Trust every other account to assume the role.

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:
Explanation:
There is a concept of Permission Boundary vs Actual IAM Policies. That is, we have a concept of “Allow” vs “Grant”. In terms of boundaries, we have the

following three boundaries:
1. SCP
2. User/Role boundaries
3. Session boundaries (ex. AssumeRole ...)
In terms of actual permission granting, we have the following:
1. Identity Policies
2. Resource Policies

QUESTION 380
A user has created a VPC with CIDR 20.0.0.0/16 using the wizard. The user has created a public subnet CIDR (20.0.0.0/24) and VPN only subnets CIDR (20.0.1.0/24) along with the VPN gateway (vgw123456) to connect to the user's data center.
The user's data center has CIDR 172.28.0.0/12. The user has also setup a NAT instance (i-123456) to allow traffic to the internet from the VPN subnet.
Which of the below mentioned options is not a valid entry for the main route table in this scenario?

A. Destination: 20.0.0.0/16 and Target: local
B. Destination: 0.0.0.0/0 and Target: i-123456
C. Destination: 172.28.0.0/12 and Target: vgw-123456
D. Destination: 20.0.1.0/24 and Target: i-123456

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The user can create subnets as per the requirement within a VPC. If the user wants to connect VPC from his own data centre, he can setup a public and VPN only subnet which uses hardware VPN access to connect with his data centre. When the user has configured this setup with Wizard, it will create a virtual private gateway to route all traffic of the VPN subnet. If the user has setup a NAT instance to route all the internet requests, then all requests to the internet should be routed to it.
All requests to the organization's DC will be routed to the VPN gateway. Here are the valid entries for the main route table in this scenario:
Destination: 0.0.0.0/0 & Target: i-123456 (To route all internet traffic to the NAT Instance) Destination: 172.28.0.0/12 & Target: vgw-123456 (To route all the organization's data centre traffic to the VPN gateway) Destination: 20.0.0.0/16 & Target: local (To allow local routing in VPC)
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario3.html

QUESTION 381
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 crossaccount 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:

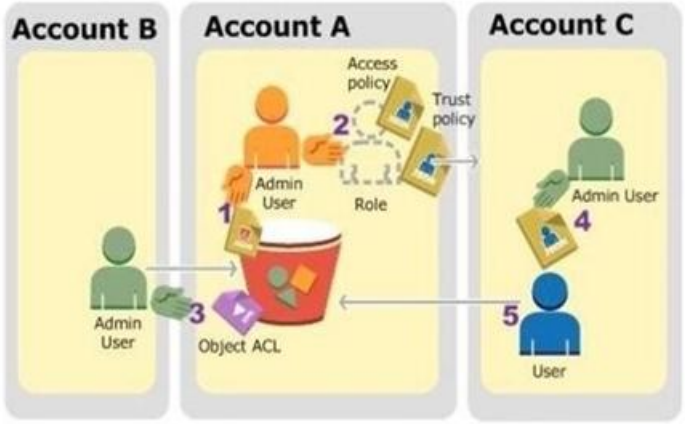

Account A administrator user attaches a bucket policy granting Account B conditional permission to upload objects.
Account A administrator creates an IAM role, establishing trust with Account C, so users in that account can access Account A. The access policy attached to the role limits what user in 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 382
A company ingests and processes streaming market data. The data rate is constant. A nightly process that calculates aggregate statistics is run, and each execution takes about 4 hours to complete. The statistical analysis is not mission critical to the business, and previous data points are picked up on the next execution if a particular run fails.

The current architecture uses a pool of Amazon EC2 Reserved Instances with 1-year reservations running full time to ingest and store the streaming data in attached Amazon EBS volumes. On-Demand EC2 instances are launched each night to perform the nightly processing, accessing the stored data from NFS shares on the ingestion servers, and terminating the nightly processing servers when complete. The Reserved Instance reservations are expiring, and the company needs to determine whether to purchase new reservations or implement a new design.
Which is the most cost-effective design?

A. Update the ingestion process to use Amazon Kinesis Data Firehose to save data to Amazon S3. Use a fleet of On- Demand EC2 instances that launches each night to perform the batch processing of the S3 data and terminates when the processing completes.
B. Update the ingestion process to use Amazon Kinesis Data Firehouse to save data to Amazon S3. Use AWS Batch to perform nightly processing with a Spot market bid of 50% of the On-Demand price.
C. Update the ingestion process to use a fleet of EC2 Reserved Instances behind a Network Load Balancer with 3-year leases. Use Batch with Spot instances with a maximum bid of 50% of the OnDemand price for the nightly processing.
D. Update the ingestion process to use Amazon Kinesis Data Firehose to save data to Amazon Redshift. Use an AWS Lambda function scheduled to run nightly with Amazon CloudWatch Events to query Amazon Redshift to generate the daily statistics.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 383
In the context of Amazon ElastiCache CLI, which of the following commands can you use to view all ElastiCache instance events for the past 24 hours?

A. elasticache-events --duration 24
B. elasticache-events --duration 1440
C. elasticache-describe-events --duration 24
D. elasticache describe-events --source-type cache-cluster --duration 1440

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In Amazon ElastiCache, the code "aws elasticache describe-events --source-type cache-cluster -- duration 1440" is used to list the cache-cluster events for the past 24 hours (1440 minutes).
Reference: http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/ECEvents.Viewing.html

QUESTION 384
A company is migrating its on-premises systems to AWS. The user environment consists of the following systems:
Windows and Linux virtual machines running on VMware. Physical servers running Red Hat Enterprise Linux.
The company wants to be able to perform the following steps before migrating to AWS:
Identify dependencies between on-premises systems.
Group systems together into applications to build migration plans.
Review performance data using Amazon Athena to ensure that Amazon EC2 instances are right-sized.
How can these requirements be met?

A. Populate the AWS Application Discovery Service import template with information from an on-premises configuration management database (CMDB). Upload the completed import template to Amazon S3, then import the data into Application Discovery Service.
B. Install the AWS Application Discovery Service Discovery Agent on each of the on-premises systems. Allow the Discovery Agent to collect data for a period of time.
C. Install the AWS Application Discovery Service Discovery Connector on each of the on-premises systems and in VMware vCenter. Allow the Discovery Connector to collect data for one week.
D. Install the AWS Application Discovery Service Discovery Agent on the physical on-premises servers. Install the AWS Application Discovery Service Discovery Connector in VMware vCenter. Allow the Discovery Agent to collect data for a period of time.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 385
If no explicit deny is found while applying IAM's Policy Evaluation Logic, the enforcement code looks for any _____ instructions that would apply to the request.

A. "cancel"
B. "suspend"
C. "allow"
D. "valid"

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

Explanation:
If an explicit deny is not found among the applicable policies for a specific request, IAM's Policy Evaluation Logic checks for any "allow" instructions to check if the request can be successfully completed.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_EvaluationLogic.html

QUESTION 386
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 5.6 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 387
Mike is appointed as Cloud Consultant in ABC.com. ABC has the following VPCs set- up in the US East Region:
A VPC with CIDR block 10.10.0.0/16, a subnet in that VPC with CIDR block 10.10.1.0/24 A VPC with CIDR block 10.40.0.0/16, a subnet in that VPC with CIDR block 10.40.1.0/24 ABC.com is trying to establish network connection between two subnets, a subnet with CIDR block 10.10.1.0/24 and another subnet with CIDR block 10.40.1.0/24.
Which one of the following solutions should Mike recommend to ABC.com?

A. Create 2 Virtual Private Gateways and configure one with each VPC.
B. Create 2 Internet Gateways, and attach one to each VPC.
C. Create a VPC Peering connection between both VPCs.
D. Create one EC2 instance in each subnet, assign Elastic IPs to both instances, and configure a set up Site-to-Site VPN connection between both EC2 instances.

Correct Answer: C

Section: (none)
Explanation

Explanation/Reference:
Explanation:
A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IP addresses. EC2 instances in either VPC can communicate with each other as if they are within the same network. You can create a VPC peering connection between your own VPCs, or with a VPC in another AWS account within a single region. AWS uses the existing infrastructure of a VPC to create a VPC peering connection; it is neither a gateway nor a VPN connection, and does not rely on a separate piece of physical hardware.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpc-peering.html

QUESTION 388
A company wants to host its website on AWS using serverless architecture design patterns for global customers. The company has outlined its requirements as follow:
The website should be responsive.
The website should offer minimal latency.
The website should be highly available.
Users should be able to authenticate through social identity providers such as Google, Facebook, and Amazon. There should be baseline DDoS protections for spikes in traffic.
How can the design requirements be met?

A. Use Amazon CloudFront with Amazon ECS for hosting the website. Use AWS Secrets Manager to provide user management and authentication functions. Use ECS Docker containers to build an API.
B. Use Amazon Route 53 latency routing with an Application Load Balancer and AWS Fargate in different regions for hosting the website. Use Amazon Cognito to provide user management and authentication functions. Use Amazon EKS containers to build an API.
C. Use Amazon CloudFront with Amazon S3 for hosting static web resources. Use Amazon Cognito to provide user management and authentication functions. Use Amazon API Gateway with AWS Lambda to build an API.
D. Use AWS Direct Connect with Amazon CloudFront and Amazon S3 for hosting static web resources. Use Amazon Cognito to provide user management authentication functions. Use AWS Lambda to build an API.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 389
Which of the following is NOT an advantage of using AWS Direct Connect?

A. AWS Direct Connect provides users access to public and private resources by using two different connections while maintaining network separation between the public and private environments.
B. AWS Direct Connect provides a more consistent network experience than Internet-based connections.
C. AWS Direct Connect makes it easy to establish a dedicated network connection from your premises to AWS.
D. AWS Direct Connect reduces your network costs.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
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 colocation environment, which in many cases can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than Internet-based connections.
By 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.
Reference: http://aws.amazon.com/directconnect/#details

QUESTION 390
A company is running an application in the AWS Cloud. The application consists of microservices that run on a fleet of Amazon EC2 instances in multiple Availability Zones behind an Application Load Balancer. The company recently added a new REST API that was implemented in Amazon API Gateway. Some of the older microservices that run on EC2 instances need to call this new API.
The company does not want the API to be accessible from the public internet and does not want proprietary data to traverse the public internet.
What should a solutions architect do to meet these requirements?

A. Create an AWS Site-to-Site VPN connection between the VPC and the API Gateway. Use API Gateway to generate a unique API key for each microservice. Configure the API methods to require the key.
B. Create an interface VPC endpoint for API Gateway, and set an endpoint policy to only allow access to the specific API.
Add a resource policy to API Gateway to only allow access from the VPC endpoint. Change the API Gateway endpoint type to private.
C. Modify the API Gateway to use IAM authentication Update the IAM policy for the IAM role that is assigned to the EC2 instances to allow access to the API Gateway Move the API Gateway into a new VPDeploy a transit gateway and connect the VPCs.
D. Create an accelerator in AWS Global Accelerator, and connect the accelerator to the API Gateway. Update the route table for all VPC subnets with a route to the created Global Accelerator endpoint IP address. Add an API key for each service to use for authentication.

Correct Answer: B
Section: (none)

Explanation

Explanation/Reference:

QUESTION 391
A company needs to create a centralized logging architecture for all of its AWS accounts. The architecture should provide near-real-time data analysis for all AWS CloudTrail logs and VPC Flow Logs across all AWS accounts. The company plans to use Amazon Elasticsearch Service (Amazon ES) to perform log analysis in the logging account.
Which strategy a solutions architect use to meet these requirements?

A. Configure CloudTrail and VPC Flow Logs in each AWS account to send data to a centralized Amazon S3 bucket in the logging account. Create and AWS Lambda function to load data from the S3 bucket to Amazon ES in the logging account.
B. Configure CloudTrail and VPC Flow Logs to send data to a log group in Amazon CloudWatch account. Configure a CloudWatch subscription filter in each AWS account to send data to Amazon Kinesis Data Firehouse in the logging account.
Load data from Kinesis Data Firehouse into Amazon ES in the logging account.
C. Configure CloudTrail and VPC Flow Logs to send data to a separate Amazon S3 bucket in each AWS account. Create an AWS Lambda function triggered by S3 events to copy the data to a centralized logging bucket. Create another Lambda function to load data from the S3 bucket to Amazon ES in the logging account.
D. Configure CloudTrail and VPC Flow Logs to send data to a log group in Amazon CloudWatch Logs in each AWS account.
Create AWS Lambda functions in each AWS accounts to subscribe to the log groups and stream the data to an Amazon S3 bucket in the logging account.
Create another Lambda function to load data from the S3 bucket to Amazon ES in the logging account.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 392
Which of the following rules must be added to a mount target security group to access Amazon Elastic File System (EFS) from an on-premises server?

A. Configure an NFS proxy between Amazon EFS and the on-premises server to route traffic.
B. Set up a Point-To-Point Tunneling Protocol Server (PPTP) to allow secure connection.
C. Permit secure traffic to the Kerberos port 88 from the on-premises server.
D. Allow inbound traffic to the Network File System (NFS) port (2049) from the on-premises server.

Correct Answer: D

Section: (none)
Explanation

Explanation/Reference:
Explanation:
By mounting an Amazon EFS file system on an on-premises server, on-premises data can be migrated into the AWS Cloud.
Any one of the mount targets in your VPC can be used as long as the subnet of the mount target is reachable by using the AWS Direct Connect connection. To access Amazon EFS from an on-premises server, a rule must be added to the mount target security group to allow inbound traffic to the NFS port (2049) from the on-premises server.
Reference: http://docs.aws.amazon.com/efs/latest/ug/how-it-works.html

QUESTION 393
An organization is planning to host a Wordpress blog as well a joomla CMS on a single instance launched with VPC. The organization wants to have separate domains for each application and assign them using Route 53. The organization may have about ten instances each with two applications as mentioned above. While launching the instance, the organization configured two separate network interfaces (primary + ENI) and wanted to have two elastic IPs for that instance. It was suggested to use a public IP from AWS instead of an elastic IP as the number of elastic IPs is restricted.
What action will you recommend to the organization?

A. I agree with the suggestion but will prefer that the organization should use separate subnets with each ENI for different public IPs.
B. I do not agree as it is required to have only an elastic IP since an instance has more than one ENI and AWS does not assign a public IP to an instance with multiple ENIs.
C. I do not agree as AWS VPC does not attach a public IP to an ENI; so the user has to use only an elastic IP only.
D. I agree with the suggestion and it is recommended to use a public IP from AWS since the organization is going to use DNS with Route 53.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. An Elastic Network Interface (ENI) is a virtual network interface that the user can attach to an instance in a VPC. The user can attach up to two ENIs with a single instance. However, AWS cannot assign a public IP when there are two ENIs attached to a single instance. It is recommended to assign an elastic IP in this scenario. If the organization wants more than 5 EIPs they can request AWS to increase the number.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html

QUESTION 394
A company plans to move regulated and security-sensitive businesses to AWS. The Security team is developing a framework to validate the adoption of AWS best practices and industry-recognized compliance standards. The AWS Management Console is the preferred method for teams to provision resources.
Which strategies should a Solutions Architect use to meet the business requirements and continuously assess, audit, and monitor the configurations of AWS

resources? (Choose two.)

A. Use AWS Config rules to periodically audit changes to AWS resources and monitor the compliance of the configuration. Develop AWS Config custom rules using AWS Lambda to establish a test-driven development approach, and further automate the evaluation of configuration changes against the required controls.
B. Use Amazon CloudWatch Logs agent to collect all the AWS SDK logs. Search the log data using a pre-defined set of filter patterns that matches mutating API calls. Send notifications using Amazon CloudWatch alarms when unintended changes are performed. Archive log data by using a batch export to Amazon S3 and then Amazon Glacier for a long-term retention and auditability.
C. Use AWS CloudTrail events to assess management activities of all AWS accounts. Ensure that CloudTrail is enabled in all accounts and available AWS services. Enable trails, encrypt CloudTrail event log files with an AWS KMS key, and monitor recorded activities with CloudWatch Logs.
D. Use the Amazon CloudWatch Events near-real-time capabilities to monitor system events patterns, and trigger AWS Lambda functions to automatically revert non-authorized changes in AWS resources. Also, target Amazon SNS topics to enable notifications and improve the response time of incident responses.
E. Use CloudTrail integration with Amazon SNS to automatically notify unauthorized API activities. Ensure that CloudTrail is enabled in all accounts and available AWS services. Evaluate the usage of Lambda functions to automatically revert nonauthorized changes in AWS resources.

Correct Answer: AC
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudwatch-alarms-for-cloudtrail.html
https://docs.aws.amazon.com/en_pv/awscloudtrail/latest/userguide/best-practices-security.html

QUESTION 395
In Amazon CloudWatch, you can publish your own metrics with the put-metric-data command. When you create a new metric using the put-metric-data command, it can take up to two minutes before you can retrieve statistics on the new metric using the get-metric-statistics command.
How long does it take before the new metric appears in the list of metrics retrieved using the list- metrics command?

A. After 2 minutes
B. Up to 15 minutes
C. More than an hour
D. Within a minute

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
You can publish your own metrics to CloudWatch with the put-metric-data command (or its Query API equivalent PutMetricData). When you create a new metric using the put-metric-data command, it can take up to two minutes before you can retrieve statistics on the new metric using the get-metric-statistics command. However, it can take up to fifteen minutes before the new metric appears in the list of metrics retrieved using the list-metrics command.
Reference: http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/publishingMetrics.html

QUESTION 396
A company manages multiple AWS accounts by using AWS Organizations. Under the root OU, the company has two OUs:
Research and DataOps.
Because of regulatory requirements, all resources that the company deploys in the organization must reside in the apnortheast- 1 Region. Additionally, EC2 instances that the company deploys in the DataOps OU must use a predefined list of instance types.
A solutions architect must implement a solution that applies these restrictions. The solution must maximize operational efficiency and must minimize ongoing maintenance.
Which combination of steps will meet these requirements? (Choose two.)

A. Create an IAM role in one account under the DataOps OU. Use the ec2:InstanceType condition key in an inline policy on the role to restrict access to specific instance type.
B. Create an IAM user in all accounts under the root OU. Use the aws:RequestedRegion condition key in an inline policy on each user to restrict access to all AWS Regions except ap-northeast-1.
C. Create an SCP. Use the aws:RequestedRegion condition key to restrict access to all AWS Regions except ap-northeast- 1. Apply the SCP to the root OU.
D. Create an SCP. Use the ec2:Region condition key to restrict access to all AWS Regions except ap-northeast-1. Apply the SCP to the root OU, the DataOps OU, and the Research OU.
E. Create an SCP. Use the ec2:InstanceType condition key to restrict access to specific instance types. Apply the SCP to the DataOps OU.

Correct Answer: BC
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_examples_aws_deny-requestedregion.html https://summitroute.com/blog/2020/03/25/aws_scp_best_practices/

QUESTION 397
In an AWS CloudFormation template, each resource declaration includes:

A. a logical ID, a resource type, and resource properties
B. a variable resource name and resource attributes

C. an IP address and resource entities
D. a physical ID, a resource file, and resource data

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In AWS CloudFormation, each resource declaration includes three parts: a logical ID that is unique within the template, a resource type, and resource properties.
Reference: http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/concept-resources.html

QUESTION 398
A Solutions Architect must create a cost-effective backup solution for a company's 500MB source code repository of proprietary and sensitive applications. The repository runs on Linux and backs up daily to tape. Tape backups are stored for 1 year.
The current solution is not meeting the company's needs because it is a manual process that is prone to error, expensive to maintain, and does not meet the need for a Recovery Point Objective (RPO) of 1 hour or Recovery Time Objective (RTO) of 2 hours. The new disaster recovery requirement is for backups to be stored offsite and to be able to restore a single file if needed.
Which solution meets the customer's needs for RTO, RPO, and disaster recovery with the LEAST effort and expense?

A. Replace local tapes with an AWS Storage Gateway virtual tape library to integrate with current backup software. Run backups nightly and store the virtual tapes on Amazon S3 standard storage in USEAST-1. Use cross-region replication to create a second copy in US-WEST-2. Use Amazon S3 lifecycle policies to perform automatic migration to Amazon Glacier and deletion of expired backups after 1 year.
B. Configure the local source code repository to synchronize files to an AWS Storage Gateway file Amazon gateway to store backup copies in an Amazon S3 Standard bucket. Enable versioning on the Amazon S3 bucket. Create Amazon S3 lifecycle policies to automatically migrate old versions of objects to Amazon S3 Standard - Infrequent Access, then Amazon Glacier, then delete backups after 1 year.
C. Replace the local source code repository storage with a Storage Gateway stored volume. Change the default snapshot frequency to 1 hour. Use Amazon S3 lifecycle policies to archive snapshots to Amazon Glacier and remove old snapshots after 1 year. Use cross-region replication to create a copy of the snapshots in US-WEST-2.
D. Replace the local source code repository storage with a Storage Gateway cached volume. Create a snapshot schedule to take hourly snapshots. Use an Amazon CloudWatch Events schedule expression rule to run an hourly AWS Lambda task to copy snapshots from US-EAST -1 to US-WEST-2.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://d1.awsstatic.com/whitepapers/aws-storage-gateway-file-gateway-for-hybrid-architectures.pdf

QUESTION 399
A company is implementing a multi-account strategy; however, the Management team has expressed concerns that services like DNS may become overly complex. The company needs a solution that allows private DNS to be shared among virtual private clouds (VPCs) in different accounts. The company will have approximately 50 accounts in total.
What solution would create the LEAST complex DNS architecture and ensure that each VPC can resolve all AWS resources?

A. Create a shared services VPC in a central account, and create a VPC peering connection from the shared services VPC to each of the VPCs in the other accounts. Within Amazon Route 53, create a privately hosted zone in the shared services VPC and resource record sets for the domain and subdomains. Programmatically associate other VPCs with the hosted zone.
B. Create a VPC peering connection among the VPCs in all accounts. Set the VPC attributes enableDnsHostnames and enableDnsSupport to “true” for each VPC. Create an Amazon Route 53 private zone for each VPC. Create resource record sets for the domain and subdomains. Programmatically associate the hosted zones in each VPC with the other VPCs.
C. Create a shared services VPC in a central account. Create a VPC peering connection from the VPCs in other accounts to the shared services VPCreate an Amazon Route 53 privately hosted zone in the shared services VPC with resource record sets for the domain and subdomains. Allow UDP and TCP port 53 over the VPC peering connections.
D. Set the VPC attributes enableDnsHostnames and enableDnsSupport to “false” in every VPC. Create an AWS Direct Connect connection with a private virtual interface. Allow UDP and TCP port 53 over the virtual interface. Use the onpremises DNS servers to resolve the IP addresses in each VPC on AWS.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 400
A company is building an electronic document management system in which users upload their documents. The application stack is entirely serverless and runs on AWS in the eu-central-1 Region. The system includes a web application that uses an Amazon CloudFront distribution for delivery with Amazon S3 as the origin. The web application communicates with Amazon API Gateway Regional endpoints. The API Gateway APIs call AWS Lambda functions that store metadata in an Amazon Aurora Serverless database and put the documents into an S3 bucket.
The company is growing steadily and has completed a proof of concept with its largest customer. The company must improve latency outside of Europe Which combination of actions will meet these requirements? (Choose two.)

A. Enable S3 Transfer Acceleration on the S3 bucket. Ensure that the web application uses the Transfer Acceleration signed URLs.
B. Create an accelerator in AWS Global Accelerator. Attach the accelerator to the CloudFront distribution.
C. Change the API Gateway Regional endpoints to edge-optimized endpoints.
D. Provision the entire stack in two other locations that are spread across the world. Use global databases on the Aurora Serverless cluster.
E. Add an Amazon RDS proxy between the Lambda functions and the Aurora Serverless database.

Correct Answer: BC
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/global-accelerator/faqs/
https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-endpoint-types.html

QUESTION 401
A solutions architect is building a web application that uses an Amazon RDS for PostgreSQL DB instance. The DB instance is expected to receive many more reads than writes. The solutions architect needs to ensure that the large amount of read traffic can be accommodated and that the DB instance is highly available.
Which steps should the solutions architect take to meet these requirements? (Choose three.)

A. Create multiple read replicas and put them into an Auto Scaling group.
B. Create multiple read replicas in different Availability Zones.
C. Create an Amazon Route 53 hosted zone and a record set for each read replica with a TTL and a weighted routing policy.
D. Create an Application Load Balancer (ALB) and put the read replicas behind the ALB.
E. Configure an Amazon CloudWatch alarm to detect a failed read replicas. Set the alarm to directly invoke an AWS Lambda function to delete its Route 53 record set.
F. Configure an Amazon Route 53 health check for each read replica using its endpoint.

Correct Answer: BDE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 402
A company is using AWS CodePipeline for the CI/CD of an application to an Amazon EC2 Auto Scaling group. All AWS resources are defined in AWS CloudFormation templates. The application artifacts are stored in an Amazon S3 bucket and deployed to the Auto Scaling group using instance user data scripts. As the application has become more complex, recent resource changes in the CloudFormation templates have caused unplanned downtime.
How should a solutions architect improve the CI/CD pipeline to reduce the likelihood that changes in the templates will cause downtime?

A. Adapt the deployment scripts to detect and report CloudFormation error conditions when performing deployments. Write test plans for a testing team to execute in a non-production environment before approving the change for production.

B. Implement automated testing using AWS CodeBuild in a test environment. Use CloudFormation change sets to evaluate changes before deployment. Use AWS CodeDeploy to leverage blue/green deployment patterns to allow evaluations and the ability to revert changes, if needed.

C. Use plugins for the integrated development environment (IDE) to check the templates for errors, and use the AWS CLI to validate that the templates are correct. Adapt the deployment code to check for error conditions and generate notifications on errors. Deploy to a test environment and execute a manual test plan before approving the change for production.

D. Use AWS CodeDeploy and a blue/green deployment pattern with CloudFormation to replace the user data deployment scripts. Have the operators log in to running instances and go through a manual test plan to verify the application is running as expected.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 403
An on-premises application will be migrated to the cloud. The application consists of a single Elasticsearch virtual machine with data source feeds from local systems that will not be migrated, and a Java web application on Apache Tomcat running on three virtual machines. The Elasticsearch server currently uses 1 TB of storage out of 16 TB available storage, and the web application is updated every 4 months. Multiple users access the web application from the Internet. There is a 10Gbit AWS Direct Connect connection established, and the application can be migrated over a scheduled 48-hour change window.
Which strategy will have the LEAST impact on the Operations staff after the migration?

A. Create an Elasticsearch server on Amazon EC2 right-sized with 2 TB of Amazon EBS and a public AWS Elastic Beanstalk environment for the web application. Pause the data sources, export the Elasticsearch index from on premises, and import into the EC2 Elasticsearch server. Move data source feeds to the new Elasticsearch server and move users to the web application.

B. Create an Amazon ES cluster for Elasticsearch and a public AWS Elastic Beanstalk environment for the web application.
Use AWS DMS to replicate Elasticsearch data. When replication has finished, move data source feeds to the new Amazon ES cluster endpoint and move users to the new web application.

C. Use the AWS SMS to replicate the virtual machines into AWS. When the migration is complete, pause the data source feeds and start the migrated Elasticsearch and web application instances. Place the web application instances behind a public Elastic Load Balancer. Move the data source feeds to the new Elasticsearch server and move users to the new web Application Load Balancer.

D. Create an Amazon ES cluster for Elasticsearch and a public AWS Elastic Beanstalk environment for the web application.
Pause the data source feeds, export the Elasticsearch index from on premises, and import into the Amazon ES cluster.
Move the data source feeds to the new Amazon ES cluster endpoint and move users to the new web application.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 404
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 IAM 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 405
A retail company needs to provide a series of data files to another company, which is its business partner. These files are saved in an Amazon S3 bucket under Account A, which belongs to the retail company. The business partner company wants one of its IAM users, User_DataProcessor, to access the files from its own AWS account (Account B).
Which combination of steps must the companies take so that User_DataProcessor can access the S3 bucket successfully?
(Choose two.)

A. Turn on the cross-origin resource sharing (CORS) feature for the S3 bucket in Account A.
B. InAccountA, set the S3 bucket policy to the following:

```
{
    "Effect": "Allow",
    "Action": [
        "s3:GetObject",
        "s3:ListBucket"
    ],
    "Resource": "arn:aws:s3:::AccountABucketName/*"
}
```

C. InAccount A, set the S3 bucket policy to the following:

```
{
    "Effect": "Allow",
    "Principal": {
        "AWS": "arn:aws:iam::AccountB:user/User_DataProcessor"
    },
    "Action": [
        "s3:GetObject",
        "s3:ListBucket"
    ],
    "Resource": [
        "arn:aws:s3:::AccountABucketName/*"
    ]
}
```


D. InAccount B, set the permissions of User_DataProcessor to the following:

```
{
    "Effect": "Allow",
    "Action": [
        "s3:GetObject",
        "s3:ListBucket"
    ],
    "Resource": "arn:aws:s3:::AccountABucketName/*"
}
```

E. InAccount B, set the permissions of User_DataProcessor to the following:

```
{
    "Effect": "Allow",
    "Principal": {
        "AWS": "arn:aws:iam::AccountB:user/User_DataProcessor"
    },
    "Action": [
        "s3:GetObject",
        "s3:ListBucket"
    ],
    "Resource": [
        "arn:aws:s3:::AccountABucketName/*"
    ]
}
```

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 406
A company is designing a new highly available web application on AWS. The application requires consistent and reliable connectivity from the application servers in AWS to a backend REST API hosted in the company's on-premises environment.
The backend connection between AWS and on-premises will be routed over an AWS Direct Connect connection through a private virtual interface. Amazon Route 53 will be used to manage private DNS records for the application to resolve the IP address on the backend REST API.
Which design would provide a reliable connection to the backend API?

A. Implement at least two backend endpoints for the backend REST API, and use Route 53 health checks to monitor the availability of each backend endpoint and perform DNS-level failover.
B. Install a second Direct Connect connection from a different network carrier and attach it to the same virtual private gateway as the first Direct Connect connection.
C. Install a second cross connect for the same Direct Connect connection from the same network carrier, and join both connections to the same link aggregation group (LAG) on the same private virtual interface.
D. Create an IPSec VPN connection routed over the public internet from the on-premises data center to AWS and attach it to the same virtual private gateway as the Direct Connect connection.

Correct Answer: B
Section: (none)

Explanation

Explanation/Reference:

QUESTION 407
After moving an E-Commerce website for a client from a dedicated server to AWS you have also set up auto scaling to perform health checks on the instances in your group and replace instances that fail these checks. Your client has come to you with his own health check system that he wants you to use as it has proved to be very useful prior to his site running on AWS.
What do you think would be an appropriate response to this given all that you know about auto scaling and CloudWatch?

A. It is not possible to implement your own health check system due to compatibility issues.
B. It is not possible to implement your own health check system. You need to use AWSs health check system.
C. It is possible to implement your own health check system and then send the instance's health information directly from your system to CloudWatch but only in the US East (N. Virginia) region.
D. It is possible to implement your own health check system and then send the instance's health information directly from your system to CloudWatch.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Auto Scaling periodically performs health checks on the instances in your group and replaces instances that fail these checks. By default, these health checks use the results of EC2 instance status checks to determine the health of an instance. If you use a load balancer with your Auto Scaling group, you can optionally choose to include the results of Elastic Load Balancing health checks.
Auto Scaling marks an instance unhealthy if the calls to the Amazon EC2 action DescribeInstanceStatus returns any other state other than running, the system status shows impaired, or the calls to Elastic Load Balancing action DescribeInstanceHealth returns OutOfService in the instance state field.
After an instance is marked unhealthy because of an Amazon EC2 or Elastic Load Balancing health check, it is scheduled for replacement.
You can customize the health check conducted by your Auto Scaling group by specifying additional checks or by having your own health check system and then sending the instance's health information directly from your system to Auto Scaling.
Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/healthcheck.html

QUESTION 408
A photo-sharing and publishing company receives 10,000 to 150,000 images daily. The company receives the images from multiple suppliers and users registered with the service. The company is moving to AWS and wants to enrich the existing metadata by adding data using Amazon Rekognition.
The following is an example of the additional data:

```
list celebrities [name of the personality] wearing [color] looking
[happy, sad] near [location example Eiffel Tower in Paris]
```

As part of the cloud migration program, the company uploaded existing image data to Amazon S3 and told users to upload images directly to Amazon S3. What should the Solutions Architect do to support these requirements?

A. Trigger AWS Lambda based on an S3 event notification to create additional metadata using Amazon Rekognition. Use Amazon DynamoDB to store the metadata and Amazon ES to create an index. Use a web front-end to provide search capabilities backed by Amazon ES.
B. Use Amazon Kinesis to stream data based on an S3 event. Use an application running in Amazon EC2 to extract metadata from the images. Then store the data on Amazon DynamoDB and Amazon CloudSearch and create an index. Use a web front-end with search capabilities backed by CloudSearch.
C. Start an Amazon SQS queue based on S3 event notifications. Then have Amazon SQS send the metadata information to Amazon DynamoDB. An application running on Amazon EC2 extracts data from Amazon Rekognition using the API and adds data to DynamoDB and Amazon ES. Use a web front-end to provide search capabilities backed by Amazon ES.
D. Trigger AWS Lambda based on an S3 event notification to create additional metadata using Amazon Rekognition. Use Amazon RDS MySQL Multi-AZ to store the metadata information and use Lambda to create an index. Use a web front-end with search capabilities backed by Lambda.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 409
How many metrics are supported by CloudWatch for Auto Scaling?

A. 7 metrics and 5 dimension
B. 5 metrics and 1 dimension
C. 1 metric and 5 dimensions
D. 8 metrics and 1 dimension

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

AWS Auto Scaling supports both detailed as well as basic monitoring of the CloudWatch metrics. Basic monitoring happens every 5 minutes, while detailed monitoring happens every minute. It supports 8 metrics and 1 dimension.
The metrics are: GroupMinSize GroupMaxSize GroupDesiredCapacity GroupInServiceInstances GroupPendingInstances GroupStandbyInstances GroupTerminatingInstances GroupTotalInstances The dimension is AutoScalingGroupName
Reference: http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/supported_services.html

QUESTION 410
A company is migrating its on-premises build artifact server to an AWS solution. The current system consists of an Apache HTTP server that serves artifacts to clients on the local network, restricted by the perimeter firewall. The artifact consumers are largely build automation scripts that download artifacts via anonymous HTTP, which the company will be unable to modify within its migration timetable.
The company decides to move the solution to Amazon S3 static website hosting. The artifact consumers will be migrated to Amazon EC2 instances located within both public and private subnets in a virtual private cloud (VPC).
Which solution will permit the artifact consumers to download artifacts without modifying the existing automation scripts?

A. Create a NAT gateway within a public subnet of the VPC. Add a default route pointing to the NAT gateway into the route table associated with the subnets containing consumers. Configure the bucket policy to allow the s3:ListBucket and s3:GetObject actions using the condition IpAddress and the condition key aws:SourceIp matching the elastic IP address of the NAT gateway.
B. Create a VPC endpoint and add it to the route table associated with subnets containing consumers. Configure the bucket policy to allow s3:ListBucket and s3:GetObject actions using the condition StringEquals and the condition key aws:sourceVpce matching the identification of the VPC endpoint.
C. Create an IAM role and instance profile for Amazon EC2 and attach it to the instances that consume build artifacts.
Configure the bucket policy to allow the s3:ListBucket and s3:GetObjects actions for the principal matching the IAM role created.
D. Create a VPC endpoint and add it to the route table associated with subnets containing consumers. Configure the bucket policy to allow s3:ListBucket and s3:GetObject actions using the condition IpAddress and the condition key aws:SourceIp matching the VPC CIDR block.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 411
In a VPC, can you modify a set of DHCP options after you create them?

A. Yes, you can modify a set of DHCP options within 48 hours after creation and there are no VPCs associated with them.
B. Yes, you can modify a set of DHCP options any time after you create them.
C. No, you can't modify a set of DHCP options after you create them.
D. Yes, you can modify a set of DHCP options within 24 hours after creation.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
After you create a set of DHCP options, you can't modify them. If you want your VPC to use a different set of DHCP options, you must create a new set and associate them with your VPC. You can also set up your VPC to use no DHCP options at all.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_DHCP_Options.html

QUESTION 412
A company has a policy that all Amazon EC2 instances that are running a database must exist within the same subnets in a shared VPC. Administrators must follow security compliance requirements and are not allowed to directly log in to the shared account. All company accounts are members of the same organization in AWS Organizations. The number of accounts will rapidly increase as the company grows.
A solutions architect uses AWS Resource Access Manager to create a resource share in the shared account.
What is the MOST operationally efficient configuration to meet these requirements?

A. Add the VPC to the resource share. Add the account IDs as principals
B. Add all subnets within the VPC to the resource share. Add the account IDs as principals
C. Add all subnets within the VPC to the resource share. Add the organization as a principal
D. Add the VPC to the resource share. Add the organization as a principal

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/blogs/networking-and-content-delivery/vpc-sharing-a-new-approach-to-multipleaccounts-and-vpc-management/

QUESTION 413
A company is developing a new service that will be accessed using TCP on a static port. A solutions architect must ensure that the service is highly available, has redundancy across Availability Zones, and is accessible using the DNS name my.service.com, which is publicly accessible. The service must use fixed address assignments so other companies can add the addresses to their allow lists.
Assuming that resources are deployed in multiple Availability Zones in a single Region, which solution will meet these requirements?

A. Create Amazon EC2 instances with an Elastic IP address for each instance. Create a Network Load Balancer (NLB) and expose the static TCP port. Register EC2 instances with the NLB. Create a new name server record set named my.service.com, and assign the Elastic IP addresses of the EC2 instances to the record set. Provide the Elastic IP addresses of the EC2 instances to the other companies to add to their allow lists.
B. Create an Amazon ECS cluster and a service definition for the application. Create and assign public IP addresses for the ECS cluster. Create a Network Load

Balancer (NLB) and expose the TCP port. Create a target group and assign the ECS cluster name to the NLCreate a new A record set named my.service.com, and assign the public IP addresses of the ECS cluster to the record set. Provide the public IP addresses of the ECS cluster to the other companies to add to their allow lists.

C. Create Amazon EC2 instances for the service. Create one Elastic IP address for each Availability Zone. Create a Network Load Balancer (NLB) and expose the assigned TCP port. Assign the Elastic IP addresses to the NLB for each Availability Zone. Create a target group and register the EC2 instances with the NLB. Create a new A (alias) record set named my.service.com, and assign the NLB DNS name to the record set.

D. Create an Amazon ECS cluster and a service definition for the application. Create and assign public IP address for each host in the cluster. Create an Application Load Balancer (ALB) and expose the static TCP port. Create a target group and assign the ECS service definition name to the ALB. Create a new CNAME record set and associate the public IP addresses to the record set. Provide the Elastic IP addresses of the Amazon EC2 instances to the other companies to add to their allow lists.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 414
A company is migrating to the cloud. It wants to evaluate the configurations of virtual machines in its existing data center environment to ensure that it can size new Amazon EC2 instances accurately. The company wants to collect metrics, such as CPU, memory, and disk utilization, and it needs an inventory of what processes are running on each instance. The company would also like to monitor network connections to map communications between servers.
Which would enable the collection of this data MOST cost effectively?

A. Use AWS Application Discovery Service and deploy the data collection agent to each virtual machine in the data center.

B. Configure the Amazon CloudWatch agent on all servers within the local environment and publish metrics to Amazon CloudWatch Logs.

C. Use AWS Application Discovery Service and enable agentless discovery in the existing virtualization environment.

D. Enable AWS Application Discovery Service in the AWS Management Console and configure the corporate firewall to allow scans over a VPN.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 415
A startup company hosts a fleet of Amazon EC2 instances in private subnets using the latest Amazon Linux 2 AMI. The company's engineers rely heavily on SSH access to the instances for troubleshooting.
The company's existing architecture includes the following:

A VPC with private and public subnets, and a NAT gateway
Site-to-Site VPN for connectivity with the on-premises environment
EC2 security groups with direct SSH access from the on-premises environment The company needs to increase security controls around SSH access and provide auditing of commands run by the engineers.
Which strategy should a solutions architect use?

A. Install and configure EC2 Instance Connect on the fleet of EC2 instances. Remove all security group rules attached to EC2 instances that allow inbound TCP on port 22. Advise the engineers to remotely access the instances by using the EC2 Instance Connect CLI.
B. Update the EC2 security groups to only allow inbound TCP on port 22 to the IP addresses of the engineer's devices.
Install the Amazon CloudWatch agent on all EC2 instances and send operating system audit logs to CloudWatch Logs.
C. Update the EC2 security groups to only allow inbound TCP on port 22 to the IP addresses of the engineer's devices.
Enable AWS Config for EC2 security group resource changes. Enable AWS Firewall Manager and apply a security group policy that automatically remediates changes to rules.
D. Create an IAM role with the AmazonSSMManagedInstanceCore managed policy attached. Attach the IAM role to all the EC2 instances. Remove all security group rules attached to the EC2 instances that allow inbound TCP on port 22. Have the engineers install the AWS Systems Manager Session Manager plugin for their devices and remotely access the instances by using the start-session API call from Systems Manager.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/Install-CloudWatch-Agent.html

QUESTION 416
A company hosts a blog post application on AWS using Amazon API Gateway, Amazon DynamoDB, and AWS Lambda. The application currently does not use API keys to authorize requests. The API model is as follows:
GET/posts/[postid] to get post details
GET/users[userid] to get user details
GET/comments/[commentid] to get comments details
The company has noticed users are actively discussing topics in the comments section, and the company wants to increase user engagement by marking the comments appears in real time.
Which design should be used to reduce comment latency and improve user experience?

A. Use edge-optimized API with Amazon CloudFront to cache API responses.
B. Modify the blog application code to request GET comment[commented] every 10 seconds.
C. Use AWS AppSync and leverage WebSockets to deliver comments.
D. Change the concurrency limit of the Lambda functions to lower the API response time.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 417
A company is running an email application across multiple AWS Regions. The company uses Ohio (us-east-2) as the primary Region and Northern Virginia (us-east-1) as the Disaster Recovery (DR) Region. The data is continuously replicated from the primary Region to the DR Region by a single instance on the public subnet in both Regions. The replication messages between the Regions have a significant backlog during certain times of the day. The backlog clears on its own after a short time, but it affects the application's RPO. Which of the following solutions should help remediate this performance problem? (Choose two.)

A. Increase the size of the instances.
B. Have the instance in the primary Region write the data to an Amazon SQS queue in the primary Region instead, and have the instance in the DR Region poll from this queue.
C. Use multiple instances on the primary and DR Regions to send and receive the replication data.
D. Change the DR Region to Oregon (us-west-2) instead of the current DR Region.
E. Attach an additional elastic network interface to each of the instances in both Regions and set up load balancing between the network interfaces.

Correct Answer: CE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 418
A company currently uses a single 1 Gbps AWS Direct Connect connection to establish connectivity between an AWS Region and its data center. The company has five Amazon VPCs, all of which are connected to the data center using the same Direct Connect connection. The Network team is worried about the single point of failure and is interested in improving the redundancy of the connections to AWS while keeping costs to a minimum.
Which solution would improve the redundancy of the connection to AWS while meeting the cost requirements?

A. Provision another 1 Gbps Direct Connect connection and create new VIFs to each of the VPCs. Configure the VIFs in a load balancing fashion using BGP.
B. Set up VPN tunnels from the data center to each VPC. Terminate each VPN tunnel at the virtual private gateway (VGW) of the respective VPC and set up BGP for route management.
C. Set up a new point-to-point Multiprotocol Label Switching (MPLS) connection to the AWS Region that's being used.
Configure BGP to use this new circuit as passive, so that no traffic flows through this unless the AWS Direct Connect fails.
D. Create a public VIF on the Direct Connect connection and set up a VPN tunnel which will terminate on the virtual private gateway (VGW) of the respective

VPC using the public VIF. Use BGP to handle the failover to the VPN connection.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 419
A Solutions Architect must design a highly available, stateless, REST service. The service will require multiple persistent storage layers for service object meta information and the delivery of content. Each request needs to be authenticated and securely processed. There is a requirement to keep costs as low as possible.
How can these requirements be met?

A. Use AWS Fargate to host a container that runs a self-contained REST service. Set up an Amazon ECS service that is fronted by an Application Load Balancer (ALB). Use a custom authenticator to control access to the API. Store request meta information in Amazon DynamoDB with Auto Scaling and static content in a secured S3 bucket. Make secure signed requests for Amazon S3 objects and proxy the data through the REST service interface.
B. Use AWS Fargate to host a container that runs a self-contained REST service. Set up an ECS service that is fronted by a cross-zone ALUse an Amazon Cognito user pool to control access to the API. Store request meta information in DynamoDB with Auto Scaling and static content in a secured S3 bucket. Generate presigned URLs when returning references to content stored in Amazon S3.
C. Set up Amazon API Gateway and create the required API resources and methods. Use an Amazon Cognito user pool to control access to the API. Configure the methods to use AWS Lambda proxy integrations, and process each resource with a unique AWS Lambda function. Store request meta information in DynamoDB with Auto Scaling and static content in a secured S3 bucket. Generate presigned URLs when returning references to content stored in Amazon S3.
D. Set up Amazon API Gateway and create the required API resources and methods. Use an Amazon API Gateway custom authorizer to control access to the API. Configure the methods to use AWS Lambda custom integrations, and process each resource with a unique Lambda function. Store request meta information in an Amazon ElastiCache Multi-AZ cluster and static content in a secured S3 bucket. Generate presigned URLs when returning references to content stored in Amazon S3.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 420
A company is deploying a distributed in-memory database on a fleet of Amazon EC2 instances. The fleet consists of a primary node and eight worker nodes. The primary node is responsible for monitoring cluster health, accepting user requests, distributing user requests to worker nodes, and sending an aggregate response back to a client. Worker nodes communicate with each other to replicate data partitions.

The company requires the lowest possible networking latency to achieve maximum performance.
Which solution will meet these requirements?

A. Launch memory optimized EC2 instances in a partition placement group.
B. Launch compute optimized EC2 instances in a partition placement group.
C. Launch memory optimized EC2 instances in a cluster placement group
D. Launch compute optimized EC2 instances in a spread placement group.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 421
A user wants to configure AutoScaling which scales up when the CPU utilization is above 70% and scales down when the CPU utilization is below 30%.
How can the user configure AutoScaling for the above mentioned condition?

A. Configure ELB to notify AutoScaling on load increase or decrease
B. Use AutoScaling with a schedule
C. Use AutoScaling by manually modifying the desired capacity during a condition
D. Use dynamic AutoScaling with a policy

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The user can configure the AutoScaling group to automatically scale up and then scale down based on the specified conditions. To configure this, the user must setup policies which will get triggered by the CloudWatch alarms.
Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/as-scale-based-on-demand.html

QUESTION 422
In DynamoDB, which of the following allows you to set alarms when you reach a specified threshold for a metric?

A. Alarm Signal
B. DynamoDB Analyzer

C. CloudWatch
D. DynamoDBALARM

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
CloudWatch allows you to set alarms when you reach a specified threshold for a metric.
Reference: http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/MonitoringDynamoDB.html

QUESTION 423
A financial company with multiple departments wants to expand its on-premises environment to the AWS Cloud. The company must retain centralized access control using an existing on-premises Active Directory (AD) service. Each department should be allowed to create AWS accounts with preconfigured networking and should have access to only a specific list of approved services. Departments are not permitted to have account administrator permissions.
What should a solutions architect do to meet these security requirements?

A. Configure AWS Identity and Access Management (IAM) with a SAML identity provider (IdP) linked to the on-premises Active Directory, and create a role to grant access. Configure AWS Organizations with SCPs and create new member accounts. Use AWS CloudFormation templates to configure the member account networking.
B. Deploy an AWS Control Tower landing zone. Create an AD Connector linked to the on-premises Active Directory. Change the identity source in AWS Single Sign-On to use Active Directory. Allow department administrators to use Account Factory to create new member accounts and networking. Grant the departments AWS power user permissions on the created accounts.
C. Deploy an Amazon Cloud Directory. Create a two-way trust relationship with the on-premises Active Directory, and create a role to grant access. Set up an AWS Service Catalog to use AWS CloudFormation templates to create the new member accounts and networking. Use IAM roles to allow access to approved AWS services.
D. Configure AWS Directory Service for Microsoft Active Directory with AWS Single Sign-On. Join the service to the onpremises Active Directory. Use AWS CloudFormation to create new member accounts and networking. Use IAM roles to allow access to approved AWS services.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://d1.awsstatic.com/whitepapers/aws-overview.pdf (46)

QUESTION 424
What is the average queue length recommended by AWS to achieve a lower latency for the 200 PIOPS EBS volume?

A. 5
B. 1
C. 2
D. 4

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The queue length is the number of pending I/O requests for a device. The optimal average queue length will vary for every customer workload, and this value depends on a particular application's sensitivity to IOPS and latency. If the workload is not delivering enough I/O requests to maintain the optimal average queue length, then the EBS volume might not consistently deliver the IOPS that have been provisioned. However, if the workload maintains an average queue length that is higher than the optimal value, then the per-request I/O latency will increase; in this case, the user should provision more IOPS for his volume. AWS recommends that the user should target an optimal average queue length of 1 for every 200 provisioned IOPS and tune that value based on his application requirements.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-workload-demand.html

QUESTION 425
A United Kingdom (UK) company recently completed a successful proof of concept in Amazon WorkSpaces. The company also has a large office in the United States (US). Staff members from each office regularly travel between the two locations and need access to a corporate WorkSpace without any reconfiguration of their WorkSpaces client.
The company has purchased a domain by using Amazon Route 53 for the connection alias. The company will use a Windows profile and document management solution.
A solutions architect needs to design the full solution. The solution must use a configuration of WorkSpaces in two AWS Regions and must provide Regional resiliency.
Which solution will meet these requirements?

A. Create a connection alias in a UK Region and a US Region. Associate the connection alias with a directory in the UK Region. Configure the DNS service for the domain in the connection alias. Configure a geolocation routing policy. Distribute the connection string to the WorkSpaces users.
B. Create a connection alias in a UK Region. Associated the connection alias with a directory in the UK Region. Configure the DNS service for the domain in the connection alias. Configure a weighted routing policy, with the UK Region set to 1 and a US Region set to 255. Distribute the connection string for the UK Region to the WorkSpaces users.
C. Create a connection alias in a UK Region and a US Region. Associate the connection aliases with a directory in each Region. Configure the DNS service for the domain in the connection alias. Configure a geolocation routing policy. Distribute the connection string to the WorkSpaces users.
D. Create a connection alias in a US Region. Associated the connection alias with a directory in the UK Region. Configure the DNS service for the domain in the connection alias. Configure a multivalue answer routing policy. Distribute the connection string for the US Region to the WorkSpaces users.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/workspaces/latest/adminguide/cross-region-redirection.html

QUESTION 426
A company wants to migrate a 30 TB Oracle data warehouse from on premises to Amazon Redshift. The company used the AWS Schema Conversion Tool (AWS SCT) to convert the schema of the existing data warehouse to an Amazon Redshift schema. The company also used a migration assessment report to identify manual tasks to complete.
The company needs to migrate the data to the new Amazon Redshift cluster during an upcoming data freeze period of 2 weeks. The only network connection between the on-premises data warehouse and AWS is a 50 Mbps internet connection.
Which migration strategy meets these requirements?

A. Create an AWS Database Migration Service (AWS DMS) replication instance. Authorize the public IP address of the replication instance to reach the data warehouse through the corporate firewall. Create a migration task to run at the beginning of the fata freeze period.
B. Install the AWS SCT extraction agents on the on-premises servers. Define the extract, upload, and copy tasks to send the data to an Amazon S3 bucket. Copy the data into the Amazon Redshift cluster. Run the tasks at the beginning of the data freeze period.
C. Install the AWS SCT extraction agents on the on-premises servers. Create a Site-to-Site VPN connection. Create an AWS Database Migration Service (AWS DMS) replication instance that is the appropriate size. Authorize the IP address of the replication instance to be able to access the on-premises data warehouse through the VPN connection.
D. Create a job in AWS Snowball Edge to import data into Amazon S3. Install AWS SCT extraction agents on the onpremises servers. Define the local and AWS Database Migration Service (AWS DMS) tasks to send the data to the Snowball Edge device. When the Snowball Edge device is returned to AWS and the data is available in Amazon S3, run the AWS DMS subtask to copy the data to Amazon Redshift.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 427
In Amazon RDS for PostgreSQL, you can provision up to 3TB storage and 30,000 IOPS per database instance. For a workload with 50% writes and 50% reads running on a cr1.8xlarge instance, you can realize over 25,000 IOPS for PostgreSQL. However, by provisioning more than this limit, you may be able to achieve:

A. higher latency and lower throughput.
B. lower latency and higher throughput.

C. higher throughput only.
D. higher latency only.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
You can provision up to 3TB storage and 30,000 IOPS per database instance. For a workload with 50% writes and 50% reads running on a cr1.8xlarge instance, you can realize over 25,000 IOPS for PostgreSQL. However, by provisioning more than this limit, you may be able to achieve lower latency and higher throughput. Your actual realized IOPS may vary from the amount you provisioned based on your database workload, instance type, and database engine choice.
Reference: https://aws.amazon.com/rds/postgresql/

QUESTION 428
A company is using Amazon Aurora MySQL for a customer relationship management (CRM) application. The application requires frequent maintenance on the database and the Amazon EC2 instances on which the application runs. For AWS Management Console access, the system administrators authenticate against AWS Identity and Access Management (IAM) using an internal identity provider. For database access, each system administrator has a user name and password that have previously been configured within the database.
A recent security audit revealed that the database passwords are not frequently rotated. The company wants to replace the passwords with temporary credentials using the company's existing AWS access controls.
Which set of options will meet the company's requirements?

A. Create a new AWS Systems Manager Parameter Store entry for each database password. Enable parameter expiration to invoke an AWS Lambda function to perform password rotation by updating the parameter value. Create an IAM policy allowing each system administrator to retrieve their current password from the Parameter Store. Use the AWS CLI to retrieve credentials when connecting to the database.
B. Create a new AWS Secrets Manager entry for each database password. Configure password rotation for each secret using an AWS Lambda function in the same VPC as the database cluster. Create an IAM policy allowing each system administrator to retrieve their current password. Use the AWS CLI to retrieve credentials when connecting to the database.
C. Enable IAM database authentication on the database. Attach an IAM policy to each system administrator's role to map the role to the database user name. Install the Amazon Aurora SSL certificate bundle to the system administrators' certificate trust store. Use the AWS CLI to generate an authentication token used when connecting to the database.
D. Enable IAM database authentication on the database. Configure the database to use the IAM identity provider to map the administrator roles to the database user. Install the Amazon Aurora SSL certificate bundle to the system administrators' certificate trust store. Use the AWS CLI to generate an authentication token used when connecting to the database.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/premiumsupport/knowledge-center/users-connect-rds-iam/

QUESTION 429
Someone is creating a VPC for their application hosting. He has created two private subnets in the same availability zone and created one subnet in a separate availability zone. He wants to make a High Availability system with an internal Elastic Load Balancer.
Which choice is true regarding internal ELBs in this scenario? (Choose two.)

A. Internal ELBs should only be launched within private subnets.
B. Amazon ELB service does not allow subnet selection; instead it will automatically select all the available subnets of the VPC.
C. Internal ELBs can support only one subnet in each availability zone.
D. An internal ELB can support all the subnets irrespective of their zones.

Correct Answer: AC
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Within this virtual private cloud, the user can launch AWS resources, such as elastic load balancers, and EC2 instances. There are two ELBs available with VPC: internet facing and internal (private) ELB. For internal servers, such as App servers the organization can create an internal load balancer in their VPC and then place back-end application instances behind the internal load balancer. The internal load balancer will route requests to the back-end application instances, which are also using private IP addresses and only accept requests from the internal load balancer. The Internal ELB supports only one subnet in each AZ and asks the user to select a subnet while configuring internal ELB.
Reference: http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/USVPC_creating_basic_lb.html

QUESTION 430
A user has created a VPC with public and private subnets using the VPC wizard. Which of the below mentioned statements is true in this scenario?

A. The user has to manually create a NAT instance
B. The Amazon VPC will automatically create a NAT instance with the micro size only
C. VPC updates the main route table used with the private subnet, and creates a custom route table with a public subnet
D. VPC updates the main route table used with a public subnet, and creates a custom route table with a private subnet

Correct Answer: C
Section: (none)

Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. If the user has created a public subnet, the instances in the public subnet can receive inbound traffic directly from the internet, whereas the instances in the private subnet cannot. If these subnets are created with Wizard, AWS will create a NAT instance of a smaller or higher size, respectively. The VPC has an implied router and the VPC wizard updates the main route table used with the private subnet, creates a custom route table and associates it with the public subnet.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario2.html

QUESTION 431
What is a possible reason you would need to edit claims issued in a SAML token?

A. The NameIdentifier claim cannot be the same as the username stored in AD.
B. Authentication fails consistently.
C. The NameIdentifier claim cannot be the same as the claim URI.
D. The NameIdentifier claim must be the same as the username stored in AD.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The two reasons you would need to edit claims issued in a SAML token are:
The NameIdentifier claim cannot be the same as the username stored in AD, and The app requires a different set of claim URIs.
Reference:
https://azure.microsoft.com/en-us/documentation/articles/active-directory-saml-claims-customization/

QUESTION 432
A user has created a VPC with a public subnet. The user has terminated all the instances which are part of the subnet.
Which of the below mentioned statements is true with respect to this scenario?

A. The subnet to which the instances were launched with will be deleted
B. When the user launches a new instance it cannot use the same subnet
C. The user cannot delete the VPC since the subnet is not deleted
D. Secondary network interfaces attached to the terminated instances may persist.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. When an instance is launched it will have a network interface attached with it. The user cannot delete the subnet until he terminates the instance and deletes the network interface. By default, network interfaces that are automatically created and attached to instances using the console are set to terminate when the instance terminates. However, network interfaces created using the command line interface aren't set to terminate when the instance terminates.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html

QUESTION 433
A team collects and routes behavioral data for an entire company. The company runs a Multi-AZ VPC environment with public subnets, private subnets, and in internet gateway. Each public subnet also contains a NAT gateway. Most of the company's applications read from and write to Amazon Kinesis Data Streams. Most of the workloads run in private subnets.
A solutions architect must review the infrastructure. The solution architect needs to reduce costs and maintain the function of the applications. The solutions architect uses Cost Explorer and notices that the cost in the EC2-Other category is consistently high. A further review shows that NatGateway-Bytes charges are increasing the cost in the EC2-Other category.
What should the solutions architect do to meet these requirements?

A. Enable VPC Flow Logs. Use Amazon Athena to analyze the logs for traffic that can be removed. Ensure that security groups are blocking traffic that is responsible for high costs.
B. Add an interface VPC endpoint for Kinesis Data Streams to the VPC. Ensure that applications have the correct IAM permissions to use the interface VPC endpoint.
C. Enable VPC Flow Logs and Amazon Detective. Review Detective findings for traffic that is not related to Kinesis Data Streams. Configure security groups to block that traffic.
D. Add an interface VPC endpoint for Kinesis Data Streams to the VPC. Ensure that the VPC endpoint policy allows traffic from the applications.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 434
A company is running a .NET three-tier web application on AWS. The team currently uses XL storage optimized instances to store and serve the website's image and video files on local instance storage. The company has encountered issues with data loss from replication and instance failures. The Solutions Architect has been asked to redesign this application to improve its reliability while keeping costs low.

Which solution will meet these requirements?

A. Set up a new Amazon EFS share, move all image and video files to this share, and then attach this new drive as a mount point to all existing servers. Create an Elastic Load Balancer with Auto Scaling general purpose instances. Enable Amazon CloudFront to the Elastic Load Balancer. Enable Cost Explorer and use AWS Trusted Advisor checks to continue monitoring the environment for future savings.
B. Implement Auto Scaling with general purpose instance types and an Elastic Load Balancer. Enable an Amazon CloudFront distribution to Amazon S3 and move images and video files to Amazon S3. Reserve general purpose instances to meet base performance requirements. Use Cost Explorer and AWS Trusted Advisor checks to continue monitoring the environment for future savings.
C. Move the entire website to Amazon S3 using the S3 website hosting feature. Remove all the web servers and have Amazon S3 communicate directly with the application servers in Amazon VPC.
D. Use AWS Elastic Beanstalk to deploy the .NET application. Move all images and video files to Amazon EFS. Create an Amazon CloudFront distribution that points to the EFS share. Reserve the m4.4xl instances needed to meet base performance requirements.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 435
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 the 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 instanceid 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 436
Can a Direct Connect link be connected directly to the Internet?

A. Yes, this can be done if you pay for it.
B. Yes, this can be done only for certain regions.
C. Yes
D. No

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
AWS Direct Connect is a network service that provides an alternative to using the Internet to utilize AWS cloud service.
Hence, a Direct Connect link cannot be connected to the Internet directly.
Reference: http://aws.amazon.com/directconnect/faqs/

QUESTION 437
Which system is used by Amazon Machine Images paravirtual (PV) virtualization during the boot process?

A. PV-BOOT
B. PV-AMI
C. PV-WORM
D. PV-GRUB

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Amazon Machine Images that use paravirtual (PV) virtualization use a system called PV-GRUB during the boot process. PVGRUB is a paravirtual boot loader that runs a patched version of GNU GRUB 0.97. When you start an instance, PV-GRUB starts the boot process and then chain loads the kernel specified by your image's menu.lst file.

Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/UserProvidedKernels.html

QUESTION 438
An organization is planning to host a web application in the AWS VPC. The organization does not want to host a database in the public cloud due to statutory requirements.
How can the organization setup in this scenario?

A. The organization should plan the app server on the public subnet and database in the organization's data center and connect them with the VPN gateway.
B. The organization should plan the app server on the public subnet and use RDS with the private subnet for a secure data operation.
C. The organization should use the public subnet for the app server and use RDS with a storage gateway to access as well as sync the data securely from the local data center.
D. The organization should plan the app server on the public subnet and database in a private subnet so it will not be in the public cloud.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. If the user wants to connect VPC from his own data centre, he can setup a public and VPN only subnet which uses hardware VPN access to connect with his data centre. When the user has configured this setup with Wizard, it will create a virtual private gateway to route all the traffic of the VPN subnet. If the virtual private gateway is attached with VPC and the user deletes the VPC from the console it will first automatically detach the gateway and only then delete the VPC.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

QUESTION 439
In the Amazon RDS Oracle DB engine, the Database Diagnostic Pack and the Database Tuning Pack are only available with _____.

A. Oracle Standard Edition
B. Oracle Express Edition
C. Oracle Enterprise Edition
D. None of these

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Reference:
https://blog.pythian.com/a-most-simple-cloud-is-amazon-rds-for-oracle-right-for-you/

QUESTION 440
Which of the following is true while using an IAM role to grant permissions to applications running on Amazon EC2 instances?

A. All applications on the instance share the same role, but different permissions.
B. All applications on the instance share multiple roles and permissions.
C. Multiple roles are assigned to an EC2 instance at a time.
D. Only one role can be assigned to an EC2 instance at a time.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Only one role can be assigned to an EC2 instance at a time, and all applications on the instance share the same role and permissions.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/role-usecase-ec2app.html

QUESTION 441
An organization is creating a VPC for their application hosting. The organization has created two private subnets in the same AZ and created one subnet in a separate zone. The organization wants to make a HA system with the internal ELB.
Which of these statements is true with respect to an internal ELB in this scenario?

A. ELB can support only one subnet in each availability zone.
B. ELB does not allow subnet selection; instead it will automatically select all the available subnets of the VPC.
C. If the user is creating an internal ELB, he should use only private subnets.
D. ELB can support all the subnets irrespective of their zones.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud.
The user has complete control over the virtual networking environment. Within this virtual private cloud, the user can launch AWS resources, such as an ELB,

and EC2 instances.
There are two ELBs available with VPC: internet facing and internal (private) ELB. For internal servers, such as App servers the organization can create an internal load balancer in their VPC and then place back-end application instances behind the internal load balancer.
The internal load balancer will route requests to the back-end application instances, which are also using private IP addresses and only accept requests from the internal load balancer. The Internal ELB supports only one subnet in each AZ and asks the user to select a subnet while configuring internal ELB.
Reference: http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/USVPC_creating_basic_lb.html

QUESTION 442
IAM Secure and Scalable is an organization which provides scalable and secure SAAS to its clients. They are planning to host a web server and App server on AWS VPC as separate tiers. The organization wants to implement the scalability by configuring Auto Scaling and load balancer with their app servers (middle tier) too.
Which of the below mentioned options suits their requirements?

A. Since ELB is internet facing, it is recommended to setup HAProxy as the Load balancer within the VPC.
B. Create an Internet facing ELB with VPC and configure all the App servers with it.
C. The user should make ELB with EC2-CLASSIC and enable SSH with it for security.
D. Create an Internal Load balancer with VPC and register all the App servers with it.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Within this virtual private cloud, the user can launch AWS resources, such as an ELB, and EC2 instances.
There are two ELBs available with VPC: internet facing and internal (private) ELB. For internal servers, such as App servers the organization can create an internal load balancer in their VPC and then place back-end application instances behind the internal load balancer. The internal load balancer will route requests to the back-end application instances, which are also using private IP addresses and only accept requests from the internal load balancer.
Reference: http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/vpc-loadbalancer-types.html

QUESTION 443
A company wants to manage the costs associated with a group of 20 applications that are infrequently used, but are still business-critical, by migrating to AWS. The applications are a mix of Java and Node.js spread across different instance clusters. The company wants to minimize costs while standardizing by using a single deployment methodology. Most of the applications are part of month-end processing routines with a small number of concurrent users, but they are occasionally run at other times. Average application memory consumption is less than 1 GB, though some applications use as much as 2.5 GB of memory during peak processing. The most important application in the group is a billing report written in Java that accesses multiple data sources and often for several hours.
Which is the MOST cost-effective solution?

A. Deploy a separate AWS Lambda function for each application. Use AWS CloudTrail logs and Amazon CloudWatch alarms to verify completion of critical jobs.
B. Deploy Amazon ECS containers on Amazon EC2 with Auto Scaling configured for memory utilization of 75%. Deploy an ECS task for each application being migrated with ECS task scaling. Monitor services and hosts by using Amazon CloudWatch.
C. Deploy AWS Elastic Beanstalk for each application with Auto Scaling to ensure that all requests have sufficient resources.
Monitor each AWS Elastic Beanstalk deployment by using CloudWatch alarms.
D. Deploy a new Amazon EC2 instance cluster that co-hosts all applications by using EC2 Auto Scaling and Application Load Balancers. Scale cluster size based on a custom metric set on instance memory utilization. Purchase 3-year Reserved Instance reservations equal to the GroupMaxSize parameter of the Auto Scaling group.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 444
A company is designing a data processing platform to process a large number of files in an Amazon S3 bucket and store the results in Amazon DynamoDB. These files will be processed once and must be retained for 1 year. The company wants to ensure that the original files and resulting data are highly available in multiple AWS Regions.
Which solution will meet these requirements?

A. Create an S3 CreateObject event notification to copy the file to Amazon Elastic Block Store (Amazon EBS). Use AWS DataSync to sync the files between EBS volumes in multiple Regions. Use an Amazon EC2 Auto Scaling group in multiple Regions to attach the EBS volumes. Process the files and store the results in a DynamoDB global table in multiple Regions.
Configure the S3 bucket with an S3 Lifecycle policy to move the files to S3 Glacier after 1 year.
B. Create an S3 CreateObject event notification to copy the file to Amazon Elastic File System (Amazon EFS). Use AWS DataSync to sync the files between EFS volumes in multiple Regions. Use an AWS Lambda function to process the EFS files and store the results in a DynamoDB global table in multiple Regions. Configure the S3 buckets with an S3 Lifecycle policy to move the files to S3 Glacier after 1 year.
C. Copy the files to an S3 bucket in another Region by using cross-Region replication. Create an S3 CreateObject event notification on the original bucket to push S3 file paths into Amazon EventBridge (Amazon CloudWatch Events). Use an AWS Lambda function to poll EventBridge (CloudWatch Events) to process each file and store the results in a DynamoDB table in each Region. Configure both S3 buckets to use the S3 Standard-Infrequent Access (S3 Standard-IA) storage class and an S3 Lifecycle policy to delete the files after 1 year.
D. Copy the files to an S3 bucket in another Region by using cross-Region replication. Create an S3 CreateObject event notification on the original bucket to execute an AWS Lambda function to process each file and store the results in a DynamoDB global table in multiple Regions. Configure both S3 buckets to use the S3 Standard-Infrequent Access (S3 Standard-IA) storage class and an S3 Lifecycle policy to delete the files after 1 year.

Correct Answer: A
Section: (none)

Explanation

Explanation/Reference:

QUESTION 445
Who is responsible for modifying the routing tables and networking ACLs in a VPC to ensure that a DB instance is reachable from other instances in the VPC?

A. AWS administrators
B. The owner of the AWS account
C. Amazon
D. The DB engine vendor

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
You are in charge of configuring the routing tables of your VPC as well as the network ACLs rules needed to make your DB instances accessible from all the instances of your VPC that need to communicate with it.
Reference: http://aws.amazon.com/rds/faqs/

QUESTION 446
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 four.)

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 447
A company runs a memory-intensive analytics application using on-demand Amazon EC2 C5 compute optimized instance.
The application is used continuously and application demand doubles during working hours. The application currently scales based on CPU usage. When scaling in occurs, a lifecycle hook is used because the instance requires 4 minutes to clean the application state before terminating.
Because users reported poor performance during working hours, scheduled scaling actions were implemented so additional instances would be added during working hours. The Solutions Architect has been asked to reduce the cost of the application.
Which solution is MOST cost-effective?

A. Use the existing launch configuration that uses C5 instances, and update the application AMI to include the Amazon CloudWatch agent. Change the Auto Scaling policies to scale based on memory utilization. Use Reserved Instances for the number of instances required after working hours, and use Spot Instances to cover the increased demand during working hours.
B. Update the existing launch configuration to use R5 instances, and update the application AMI to include SSM Agent.
Change the Auto Scaling policies to scale based on memory utilization. Use Reserved Instances for the number of instances required after working hours, and use Spot Instances with on-Demand instances to cover the increased demand during working hours.
C. Use the existing launch configuration that uses C5 instances, and update the application AMI to include SSM Agent.
Leave the Auto Scaling policies to scale based on CPU utilization. Use scheduled Reserved Instances for the number of instances required after working hours, and use Spot Instances to cover the increased demand during working hours.
D. Create a new launch configuration using R5 instances, and update the application AMI to include the Amazon CloudWatch agent. Change the Auto Scaling policies to scale based on memory utilization. Use Reserved Instances for the number of instances required after working hours, and use Standard Reserved Instances with On-Demand Instances to cover the increased demand during working hours.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/monitoring_ec2.html

QUESTION 448
In DynamoDB, to get a detailed listing of secondary indexes on a table, you can use the _____ action.

A. BatchGetItem
B. TableName
C. DescribeTable

D. GetItem

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In DynamoDB, DescribeTable returns information about the table, including the current status of the table, when it was created, the primary key schema, and any indexes on the table.
Reference: http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/SecondaryIndexes.html

QUESTION 449
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 450

Is there any way to own a direct connection to Amazon Web Services?

A. No, AWS only allows access from the public Internet.
B. No, you can create an encrypted tunnel to VPC, but you cannot own the connection.
C. Yes, you can via Amazon Dedicated Connection
D. Yes, you can via AWS Direct Connect.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
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.
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.html

QUESTION 451
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 452
A company's service for video game recommendations has just gone viral. The company has new users from all over the world. The website for the service is hosted on a set of Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer (ALB). The website consists of static content with different resources being loaded depending on the device type.
Users recently reported that the load time for the website has increased. Administrators are reporting high loads on the EC2 instances that host the service.
Which set actions should a solutions architect take to improve response times?

A. Create separate Auto Scaling groups based on device types. Switch to Network Load Balancer (NLB). Use the User- Agent HTTP header in the NLB to route to a different set of EC2 instances.
B. Move content to Amazon S3. Create an Amazon CloudFront distribution to serve content out of the S3 bucket. Use Lambda@Edge to load different resources based on the User-Agent HTTP header.
C. Create a separate ALB for each device type. Create one Auto Scaling group behind each ALB. Use Amazon Route 53 to route to different ALBs depending on the User-Agent HTTP header.
D. Move content to Amazon S3. Create an Amazon CloudFront distribution to serve content out of the S3 bucket. Use the User-Agent HTTP header to load different content.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 453
A company had a tight deadline to migrate its on-premises environment to AWS. It moved over Microsoft SQL Servers and Microsoft Windows Servers using the virtual machine import/export service and rebuild other applications native to the cloud.
The team created both Amazon EC2 databases and used Amazon RDS. Each team in the company was responsible for migrating their applications, and they have created individual accounts for isolation of resources. The company did not have much time to consider costs, but now it would like suggestions on reducing its AWS spend.
Which steps should a Solutions Architect take to reduce costs?

A. Enable AWS Business Support and review AWS Trusted Advisor's cost checks. Create Amazon EC2 Auto Scaling groups for applications that experience fluctuating demand. Save AWS Simple Monthly Calculator reports in Amazon S3 for trend analysis. Create a master account under Organizations and have teams join for consolidated billing.
B. Enable Cost Explorer and AWS Business Support. Reserve Amazon EC2 and Amazon RDS DB instances. Use Amazon CloudWatch and AWS Trusted Advisor for monitoring and to receive costsavings suggestions. Create a master account under Organizations and have teams join for consolidated billing.
C. Create an AWS Lambda function that changes the instance size based on Amazon CloudWatch alarms. Reserve instances based on AWS Simple Monthly Calculator suggestions. Have an AWS WellArchitected framework review and apply recommendations. Create a master account under Organizations and have teams join for consolidated billing.
D. Create a budget and monitor for costs exceeding the budget. Create Amazon EC2 Auto Scaling groups for applications that experience fluctuating demand. Create an AWS Lambda function that changes instance sizes based on Amazon CloudWatch alarms. Have each team upload their bill to an Amazon S3 bucket for analysis of team spending. Use Spot Instances on nightly batch processing jobs.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 454
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 455
A software as a service (SaaS) company offers a cloud solution for document management to private law firms and the public sector. A local government client recently mandated that highly confidential documents cannot be stored outside the country. The company CIO asks a Solutions Architect to ensure the application can adapt to this new requirement. The CIO also wants to have a proper backup plan for these documents, as backups are not currently performed.
What solution meets these requirements?

A. Tag documents that are not highly confidential as regular in Amazon S3. Create individual S3 buckets for each user. Upload objects to each user's bucket. Set S3 bucket replication from these buckets to a central S3 bucket in a different AWS account and AWS Region. Configure an AWS Lambda function triggered by scheduled events in Amazon CloudWatch to delete objects that are tagged as secret in the S3 backup bucket.

B. Tag documents as either regular or secret in Amazon S3. Create an individual S3 backup bucket in the same AWS account and AWS Region. Create a cross-region S3 bucket in a separate AWS account. Set proper IAM roles to allow crossregion permissions to the S3 buckets. Configure an AWS Lambda function triggered by Amazon CloudWatch scheduled events to copy objects that are tagged as secret to the S3 backup bucket and objects tagged as normal to the cross-region S3 bucket.

C. Tag documents as either regular or secret in Amazon S3. Create an individual S3 backup bucket in the same AWS account and AWS Region. Use S3 selective cross-region replication based on object tags to move regular documents to an S3 bucket in a different AWS Region. Configure an AWS Lambda function that triggers when new S3 objects are created in the main bucket to replicate only documents tagged as secret into the S3 bucket in the same AWS Region.

D. Tag highly confidential documents as secret in Amazon S3. Create an individual S3 backup bucket in the same AWS account and AWS Region. Use S3 selective cross-region replication based on object tags to move regular documents to a different AWS Region. Create an Amazon CloudWatch Events rule for new S3 objects tagged as secret to trigger an AWS Lambda function to replicate them into a separate bucket in the same AWS Region.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 456
Which of the following is NOT a true statement about Auto Scaling?

A. Auto Scaling can launch instances in different Azs.

B. Auto Scaling can work with CloudWatch.

C. Auto Scaling can launch an instance at a specific time.

D. Auto Scaling can launch instances in different regions.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Auto Scaling provides an option to scale up and scale down based on certain conditions or triggers from Cloudwatch. A user can configure such that Auto Scaling launches instances across Azs, but it cannot span across regions.

Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/as-dg.pdf

QUESTION 457
A user is configuring MySQL RDS with PIOPS. What should be the minimum size of DB storage provided by the user?

A. 1 TB
B. 50 GB
C. 5 GB
D. 100 GB

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
If the user is trying to enable PIOPS with MySQL RDS, the minimum size of storage should be 100 GB.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PIOPS.html

QUESTION 458
A data analytics company has an Amazon Redshift cluster that consists of several reserved nodes. The cluster is experiencing unexpected bursts of usage because a team of employees is compiling a deep audit analysis report. The queries to generate the report are complex read queries and are CPU intensive. Business requirements dictate that the cluster must be able to service read and write queries at all times. A solutions architect must devise a solution that accommodates the bursts of usage.
Which solution meets these requirements MOST cost-effectively?

A. Provision an Amazon EMR cluster. Offload the complex data processing tasks.
B. Deploy an AWS Lambda function to add capacity to the Amazon Redshift cluster by using a classic resize operation when the cluster's CPU metrics in Amazon CloudWatch reach 80%.
C. Deploy an AWS Lambda function to add capacity to the Amazon Redshift cluster by using an elastic resize operation when the cluster's CPU metrics in Amazon CloudWatch reach 80%
D. Turn on the Concurrency Scaling feature for the Amazon Redshift cluster.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 459
A Solutions Architect is redesigning an image-viewing and messaging platform to be delivered as SaaS. Currently, there is a farm of virtual desktop infrastructure (VDI) that runs a desktop image-viewing application and a desktop messaging application. Both applications use a shared database to manage user accounts and sharing. Users log in from a web portal that launches the applications and streams the view of the application on the user's machine. The Development Operations team wants to move away from using VDI and wants to rewrite the application.
What is the MOST cost-effective architecture that offers both security and ease of management?

A. Run a website from an Amazon S3 bucket with a separate S3 bucket for images and messaging data. Call AWS Lambda functions from embedded JavaScript to manage the dynamic content, and use Amazon Cognito for user and sharing management.
B. Run a website from Amazon EC2 Linux servers, storing the images in Amazon S3, and use Amazon Cognito for user accounts and sharing. Create AWS CloudFormation templates to launch the application by using EC2 user data to install and configure the application.
C. Run a website as an AWS Elastic Beanstalk application, storing the images in Amazon S3, and using an Amazon RDS database for user accounts and sharing. Create AWS CloudFormation templates to launch the application and perform blue/green deployments.
D. Run a website from an Amazon S3 bucket that authorizes Amazon AppStream to stream applications for a combined image viewer and messenger that stores images in Amazon S3. Have the website use an Amazon RDS database for user accounts and sharing.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 460
A Company had a security event whereby an Amazon S3 bucket with sensitive information was made public. Company policy is to never have public S3 objects, and the Compliance team must be informed immediately when any public objects are identified.
How can the presence of a public S3 object be detected, set to trigger alarm notifications, and automatically remediated in the future? (Choose two.)

A. Turn on object-level logging for Amazon S3. Turn on Amazon S3 event notifications to notify by using an Amazon SNS topic when a PutObject API call is made with a public-read permission.
B. Configure an Amazon CloudWatch Events rule that invokes an AWS Lambda function to secure the S3 bucket.
C. Use the S3 bucket permissions for AWS Trusted Advisor and configure a CloudWatch event to notify by using Amazon SNS.
D. Turn on object-level logging for Amazon S3. Configure a CloudWatch event to notify by using an SNS topic when a PutObject API call with public-read permission is detected in the AWS CloudTrail logs.
E. Schedule a recursive Lambda function to regularly change all object permissions inside the S3 bucket.

Correct Answer: BD
Section: (none)

Explanation

Explanation/Reference:

QUESTION 461
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.
Question was explicitly focused on security, so IAM with RDS is the best choice.
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:
Explanation:
We can use either RDS or DynamoDB, however in our given answers, IAM role is mentioned only with RDS, so I would go with Answer
B. Question was explicitly focused on security, so IAM with RDS is the best choice.

QUESTION 462
A Solutions Architect has created an AWS CloudFormation template for a three-tier application that contains an Auto Scaling group of Amazon EC2 instances running a custom AMI.
The Solutions Architect wants to ensure that future updates to the custom AMI can be deployed to a running stack by first updating the template to refer to the new AMI, and then invoking UpdateStack to replace the EC2 instances with instances launched from the new AMI.
How can updates to the AMI be deployed to meet these requirements?

A. Create a change set for a new version of the template, view the changes to the running EC2 instances to ensure that the AMI is correctly updated, and then execute the change set.
B. Edit the AWS::AutoScaling::LaunchConfiguration resource in the template, changing its DeletionPolicy to Replace.
C. Edit the AWS::AutoScaling::AutoScalingGroup resource in the template, inserting an UpdatePolicy attribute.
D. Create a new stack from the updated template. Once it is successfully deployed, modify the DNS records to point to the new stack and delete the old stack.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-updatepolicy.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-properties-as-launchconfig.html

QUESTION 463
A company wants to provide desktop as a service (DaaS) to a number of employees using Amazon WorkSpaces.
WorkSpaces will need to access files and services hosted on premises with authorization based on the company's Active Directory. Network connectivity will be provided through an existing AWS Direct Connect connection.
The solution has the following requirements:
Credentials from Active Directory should be used to access on-premises files and services.
Credentials from Active Directory should not be stored outside the company.
End users should have single sign-on (SSO) to on-premises files and services once connected to WorkSpaces.
Which strategy should the solutions architect use for end user authentication?

A. Create an AWS Directory Service for Microsoft Active Directory (AWS Managed Microsoft AD) directory within the WorkSpaces VPC. Use the Active Directory Migration Tool (ADMT) with the Password Export Server to copy users from the on-premises Active Directory to AWS Managed Microsoft AD. Set up a one-way trust allowing users from AWS Managed Microsoft AD to access resources in the on-premises Active Directory. Use AWS Managed Microsoft AD as the directory for WorkSpaces.
B. Create a service account in the on-premises Active Directory with the required permissions. Create an AD Connector in AWS Directory Service to be deployed on premises using the service account to communicate with the on-premises Active Directory. Ensure the required TCP ports are open from the WorkSpaces VPC to the on-premises AD Connector. Use the AD Connector as the directory for WorkSpaces.
C. Create a service account in the on-premises Active Directory with the required permissions. Create an AD Connector in AWS Directory Service within the WorkSpaces VPC using the service account to communicate with the on-premises Active Directory. Use the AD Connector as the directory for WorkSpaces.
D. Create an AWS Directory Service for Microsoft Active Directory (AWS Managed Microsoft AD) directory in the AWS Directory Service within the WorkSpaces VPC. Set up a one-way trust allowing users from the on-premises Active Directory to access resources in the AWS Managed Microsoft AUse AWS Managed Microsoft AD as the directory for WorkSpaces.
Create an identity provider with AWS Identity and Access Management (IAM) from an on-premises ADFS server. Allow users from this identity provider to assume a role with a policy allowing them to run WorkSpaces.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/directoryservice/latest/admin-guide/directory_microsoft_ad.html

QUESTION 464
Can Provisioned IOPS be used on RDS instances launched in a VPC?

A. Yes, they can be used only with Oracle based instances.
B. Yes, they can be used for all RDS instances.
C. No
D. Yes, they can be used only with MySQL based instances.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The basic building block of Amazon RDS is the DB instance. DB instance storage comes in three types:
Magnetic, General Purpose (SSD), and Provisioned IOPS (SSD). When you buy a server, you get CPU, memory, storage, and IOPS, all bundled together. With Amazon RDS, these are split apart so that you can scale them independently. So, for example, if you need more CPU, less IOPS, or more storage, you can easily allocate them.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/RDSFAQ.PIOPS.html

QUESTION 465
During an audit, a security team discovered that a development team was putting IAM user secret access keys in their code and then committing it to an AWS CodeCommit repository. The security team wants to automatically find and remediate instances of this security vulnerability.
Which solution will ensure that the credentials are appropriately secured automatically?

A. Run a script nightly using AWS Systems Manager Run Command to search for credentials on the development instances. If found, use AWS Secrets Manager to rotate the credentials.
B. Use a scheduled AWS Lambda function to download and scan the application code from CodeCommit. If credentials are found, generate new credentials and store them in AWS KMS.
C. Configure Amazon Macie to scan for credentials in CodeCommit repositories. If credentials are found, trigger an AWS Lambda function to disable the credentials and notify the user.

D. Configure a CodeCommit trigger to invoke an AWS Lambda function to scan new code submissions for credentials. If credentials are found, disable them in AWS IAM and notify the user.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://aws.amazon.com/blogs/security/how-to-find-update-access-keys-password-mfa-aws-management-console/

QUESTION 466
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 467
A company runs an IoT platform on AWS. IoT sensors in various locations send data to the company's Node.js API servers on Amazon EC2 instances running behind an Application Load Balancer. The data is stored in an Amazon RDS MySQL DB instance that uses a 4 TB General Purpose SSD volume.
The number of sensors the company has deployed in the field has increased over time, and is expected to grow significantly.
The API servers are consistently overloaded and RDS metrics show high write latency.
Which of the following steps together will resolve the issues permanently and enable growth as new sensors are provisioned, while keeping this platform cost-efficient? (Choose two.)

A. Resize the MySQL General Purpose SSD storage to 6 TB to improve the volume's IOPS
B. Re-architect the database tier to use Amazon Aurora instead of an RDS MySQL DB instance and add read replicas

C. Leverage Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data
D. Use AWS-X-Ray to analyze and debug application issues and add more API servers to match the load
E. Re-architect the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance

Correct Answer: CE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 468
A company deployed a three-tier web application in two regions: us-east-1 and eu-west-1. The application must be active in both regions at the same time. The database tier of the application uses a single Amazon RDS Aurora database globally, with a master in us-east-1 and a read replica in eu-west-1. Both regions are connected by a VPN.
The company wants to ensure that the application remains available even in the event of a region-level failure of all of the application's components. It is acceptable for the application to be in read-only mode for up to 1 hour. The company plans to configure two Amazon Route 53 record sets, one for each of the regions.
How should the company complete the configuration to meet its requirements while providing the lowest latency for the application end-users? (Choose two.)

A. Use failover routing and configure the us-east-1 record set as primary and the eu-west-1 record set as secondary. Configure an HTTP health check for the web application in us-east-1, and associate it to the us-east-1 record set.
B. Use weighted routing and configure each record set with a weight of 50. Configure an HTTP health check for each region,and attach it to the record set for that region.
C. Use latency-based routing for both record sets. Configure a health check for each region and attach it to the record set for that region.
D. Configure an Amazon CloudWatch alarm for the health checks in us-east-1, and have it invoke an AWS Lambda function that promotes the read replica in eu-west-1.
E. Configure Amazon RDS event notifications to react to the failure of the database in us-east-1 by invoking an AWS Lambda function that promotes the read replica in eu-west-1.

Correct Answer: CE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 469
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: E
Section: (none)
Explanation

Explanation/Reference:

QUESTION 470
A company is running a large application on premises. Its technology stack consists of Microsoft .NET for the web server platform and Apache Cassandra for the database. The company wants to migrate this application to AWS to improve service reliability. The IT team also wants to reduce the time it spends on capacity management and maintenance of this infrastructure. The Development team is willing and available to make code changes to support the migration.
Which design is the LEAST complex to manage after the migration?

A. Migrate the web servers to Amazon EC2 instances in an Auto Scaling group that is running .NET. Migrate the existing Cassandra database to Amazon Aurora with multiple read replicas, and run both in a Multi-AZ mode.
B. Migrate the web servers to an AWS Elastic Beanstalk environment that is running the .NET platform in a Multi-AZ Auto Scaling configuration. Migrate the Cassandra database to Amazon EC2 instances that are running in a Multi-AZ configuration.
C. Migrate the web servers to an AWS Elastic Beanstalk environment that is running the .NET platform in a Multi-AZ Auto Scaling configuration. Migrate the existing Cassandra database to Amazon DynamoDB.
D. Migrate the web servers to Amazon EC2 instances in an Auto Scaling group that is running .NET. Migrate the existing Cassandra database to Amazon DynamoDB.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 471
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 472
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 (NetworkIn). 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 10gbps.
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 473
Within an IAM policy, can you add an IfExists condition at the end of a Null condition?

A. Yes, you can add an IfExists condition at the end of a Null condition but not in all Regions.
B. Yes, you can add an IfExists condition at the end of a Null condition depending on the condition.
C. No, you cannot add an IfExists condition at the end of a Null condition.
D. Yes, you can add an IfExists condition at the end of a Null condition.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Within an IAM policy, IfExists can be added to the end of any condition operator except the Null condition. It can be used to indicate that conditional comparison needs to happen if the policy key is present in the context of a request; otherwise, it can be ignored.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements.html

QUESTION 474
A company stores sales transaction data in Amazon DynamoDB tables. To detect anomalous behaviors and respond quickly, all changes to the items stored in the DynamoDB tables must be logged within 30 minutes.
Which solution meets the requirements?

A. Copy the DynamoDB tables into Apache Hive tables on Amazon EMR every hour and analyze them for anomalous behaviors. Send Amazon SNS notifications when anomalous behaviors are detected.
B. Use AWS CloudTrail to capture all the APIs that change the DynamoDB tables. Send SNS notifications when anomalous behaviors are detected using CloudTrail event filtering.
C. Use Amazon DynamoDB Streams to capture and send updates to AWS Lambda. Create a Lambda function to output records to Amazon Kinesis Data Streams. Analyze any anomalies with Amazon Kinesis Data Analytics. Send SNS notifications when anomalous behaviors are detected.
D. Use event patterns in Amazon CloudWatch Events to capture DynamoDB API call events with an AWS Lambda function as a target to analyze behavior. Send SNS notifications when anomalous behaviors are detected.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 475
A Solutions Architect is designing a deployment strategy for an application tier and has the following requirements:
The application code will need a 500 GB static dataset to be present before application startup.
The application tier must be able to scale up and down based on demand with as little startup time as possible.
The Development team should be able to update the code multiple times each day.
Critical operating system (OS) patches must be installed within 48 hours of being released.
Which deployment strategy meets these requirements?

A. Use AWS Systems Manager to create a new AMI with the updated OS patches. Update the Auto Scaling group to use the patched AMI and replace existing unpatched instances. Use AWS CodeDeploy to push the application code to the instances.
Store the static data in Amazon EFS.
B. Use AWS Systems Manager to create a new AMI with updated OS patches. Update the Auto Scaling group to use the patched AMI and replace existing unpatched instances. Update the OS patches and the application code as batch job every night. Store the static data in Amazon EFS.
C. Use an Amazon-provided AMI for the OS. Configure an Auto Scaling group set to a static instance count. Configure an Amazon EC2 user data script to download the data from Amazon S3. Install OS patches with AWS Systems Manager when they are released. Use AWS CodeDeploy to push the application code to the instances.
D. Use an Amazon-provided AMI for the OS. Configure an Auto Scaling group. Configure an Amazon EC2 user data script to download the data from Amazon S3. Replace existing instances after each updated Amazon-provided AMI release. Use AWS CodeDeploy to push the application code to the instances.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 476
A company wants to use Amazon S3 to back up its on-premises file storage solution. The company's on-premises file storage solution supports NFS, and the company wants its new solution to support NFS. The company wants to archive the backup files after 5 days. If the company needs archived files for disaster recovery, the company is willing to wait a few days for the retrieval of those files.
Which solution meets these requirements MOST cost-effectively?

A. Deploy an AWS Storage Gateway files gateway that is associated with an S3 bucket. Move the files from the on-premises file storage solution to the file gateway. Create an S3 Lifecycle rule to move the file to S3 Standard-Infrequent Access (S3 Standard-IA) after 5 days.
B. Deploy an AWS Storage Gateway volume gateway that is associated with an S3 bucket. Move the files from the onpremises file storage solution to the volume gateway. Create an S3 Lifecycle rule to move the files to S3 Glacier Deep Archive after 5 days.
C. Deploy an AWS Storage Gateway tape gateway that is associated with an S3 bucket. Move the files from the on-premises file storage solution to the tape gateway. Create an S3 Lifecycle rule to move the files to S3 Standard-Infrequent Access (S3 Standard-IA) after 5 days.
D. Deploy an AWS Storage Gateway file gateway that is associated with an S3 bucket. Move the files from the on-premises file storage solution to the tape gateway. Create an S3 Lifecycle rule to move the files to S3 Standard-Infrequent Access (S3 Standard-IA) after 5 days.

E. Deploy an AWS Storage Gateway file gateway that is associated with an S3 bucket. Move the files from the on-premises file storage solution to the file gateway. Create an S3 Lifecycle rule to move the files to S3 Glacier Deep Archive after 5 days.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/blogs/database/storing-sql-server-backups-in-amazon-s3-using-aws-storage-gateway/

QUESTION 477
A retail company has a custom .NET web application running on AWS that uses Microsoft SQL Server for the database. The application servers maintain a user's session locally.
Which combination of architecture changes are needed to ensure all tiers of the solution are highly available? (Choose three.)

A. Refactor the application to store the user's session in Amazon ElastiCache. Use Application Load Balancers to distribute the load between application instances.
B. Set up the database to generate hourly snapshots using Amazon EBS. Configure an Amazon CloudWatch Events rule to launch a new database instance if the primary one fails.
C. Migrate the database to Amazon RDS for SQL Server. Configure the RDS instance to use a Multi-AZ deployment.
D. Move the .NET content to an Amazon S3 bucket. Configure the bucket for static website hosting.
E. Put the application instances in an Auto Scaling group. Configure the Auto Scaling group to create new instances if an instance becomes unhealthy.
F. Deploy Amazon CloudFront in front of the application tier. Configure CloudFront to serve content from healthy application instances only.

Correct Answer: BDE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 478
A company hosts a legacy application that runs on an Amazon EC2 instance inside a VPC without internet access. Users access the application with a desktop program installed on their corporate laptops. Communication between the laptops and the VPC flows through AWS Direct Connect (DX). A new requirement states that all data in transit must be encrypted between users and the VPC.
Which strategy should a solutions architect use to maintain consistent network performance while meeting this new requirement?

A. Create a client VPN endpoint and configure the laptops to use an AWS client VPN to connect to the VPC over the internet.
B. Create a new public virtual interface for the existing DX connection, and create a new VPN that connects to the VPC over the DX public virtual interface.
C. Create a new Site-to-Site VPN that connects to the VPC over the internet.
D. Create a new private virtual interface for the existing DX connection, and create a new VPN that connects to the VPC over the DX private virtual interface.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 479
Which of the following cache engines does Amazon ElastiCache support?

A. Amazon ElastiCache supports Memcached and Redis.
B. Amazon ElastiCache supports Redis and WinCache.
C. Amazon ElastiCache supports Memcached and Hazelcast.
D. Amazon ElastiCache supports Memcached only.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The cache engines supported by Amazon ElastiCache are Memcached and Redis.
Reference: http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/SelectEngine.html

QUESTION 480
A financial services company is moving to AWS and wants to enable developers to experiment and innovate while preventing access to production applications.
The company has the following requirements:
Production workloads cannot be directly connected to the internet.
All workloads must be restricted to the us-west-2 and eu-central-1 Regions.
Notification should be sent when developer sandboxes exceed $500 in AWS spending monthly.
Which combination of actions needs to be taken to create a multi-account structure that meets the company's requirements?
(Choose three.)

A. Create accounts for each production workload within an organization in AWS Organizations. Place the production accounts within an organizational unit (OU).

For each account, delete the default VPC. Create an SCP with a Deny rule for the attach an internet gateway and create a default VPC actions. Attach the SCP to the OU for the production accounts.

B. Create accounts for each production workload within an organization in AWS Organizations. Place the production accounts within an organizational unit (OU). Create an SCP with a Deny rule on the attach an internet gateway action.
Create an SCP with a Deny rule to prevent use of the default VPC. Attach the SCPs to the OU for the production accounts.

C. Create a SCP containing a Deny Effect for cloudfront:*, iam:*, route53:*, and support:* with a StringNotEquals condition on an aws:RequestedRegion condition key with us-west-2 and eu-central-1 values. Attach the SCP to the organization's root.

D. Create an IAM permission boundary containing a Deny Effect for cloudfront:*, iam:*, route53:*, and support:* with a StringNotEquals condition on an aws:RequestedRegion condition key with us-west-2 and eu-central-1 values. Attach the permission boundary to an IAM group containing the development and production users.

E. Create accounts for each development workload within an organization in AWS Organizations. Place the development accounts within an organizational unit (OU). Create a custom AWS Config rule to deactivate all IAM users when an account's monthly bill exceeds $500.

F. Create accounts for each development workload within an organization in AWS Organizations. Place the development accounts within an organizational unit (OU). Create a budget within AWS Budgets for each development account to monitor and report on monthly spending exceeding $500.

Correct Answer: ACF
Section: (none)
Explanation

Explanation/Reference:

QUESTION 481
A company's application is increasingly popular and experiencing latency because of high volume reads on the database server.
The service has the following properties:
A highly available REST API hosted in one region using Application Load Balancer (ALB) with auto scaling. A MySQL database hosted on an Amazon EC2 instance in a single Availability Zone.
The company wants to reduce latency, increase in-region database read performance, and have multi-region disaster recovery capabilities that can perform a live recovery automatically without any data or performance loss (HA/DR).
Which deployment strategy will meet these requirements?

A. Use AWS CloudFormation StackSets to deploy the API layer in two regions. Migrate the database to an Amazon Aurora with MySQL database cluster with multiple read replicas in one region and a read replica in a different region than the source database cluster. Use Amazon Route 53 health checks to trigger a DNS failover to the standby region if the health checks to the primary load balancer fail. In the event of Route 53 failover, promote the cross-region database replica to be the master and build out new read replicas in the standby region.

B. Use Amazon ElastiCache for Redis Multi-AZ with an automatic failover to cache the database read queries. Use AWS OpsWorks to deploy the API layer, cache layer, and existing database layer in two regions. In the event of failure, use Amazon Route 53 health checks on the database to trigger a DNS failover to the standby region if the health checks in the primary region fail. Back up the MySQL database frequently, and in the event of a failure in an active region, copy the backup to the standby region and restore the standby database.

C. Use AWS CloudFormation StackSets to deploy the API layer in two regions. Add the database to an Auto Scaling group. Add a read replica to the database in the second region. Use Amazon Route 53 health checks on the database to trigger a DNS failover to the standby region if the health checks in the primary region fail. Promote the cross-region database replica to be the master and build out new read replicas in the standby region.

D. Use Amazon ElastiCache for Redis Multi-AZ with an automatic failover to cache the database read queries. Use AWS OpsWorks to deploy the API layer, cache layer, and existing database layer in two regions. Use Amazon Route 53 health checks on the ALB to trigger a DNS failover to the standby region if the health checks in the primary region fail. Back up the MySQL database frequently, and in the event of a failure in an active region, copy the backup to the standby region and restore the standby database.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 482
In Amazon ElastiCache, which of the following statements is correct?

A. When you launch an ElastiCache cluster into an Amazon VPC private subnet, every cache node is assigned a public IP address within that subnet.
B. You cannot use ElastiCache in a VPC that is configured for dedicated instance tenancy.
C. If your AWS account supports only the EC2-VPC platform, ElastiCache will never launch your cluster in a VPC.
D. ElastiCache is not fully integrated with Amazon Virtual Private Cloud (VPC).

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The VPC must allow non-dedicated EC2 instances. You cannot use ElastiCache in a VPC that is configured for dedicated instance tenancy.
Reference: http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/AmazonVPC.EC.html

QUESTION 483
A company has an application written using an in-house software framework. The framework installation takes 30 minutes and is performed with a user data script. Company Developers deploy changes to the application frequently. The framework installation is becoming a bottleneck in this process.
Which of the following would speed up this process?

A. Create a pipeline to build a custom AMI with the framework installed and use this AMI as a baseline for application deployments.

B. Employ a user data script to install the framework but compress the installation files to make them smaller.
C. Create a pipeline to parallelize the installation tasks and call this pipeline from a user data script.
D. Configure an AWS OpsWorks cookbook that installs the framework instead of employing user data. Use this cookbook as a base for all deployments.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/codepipeline/features/?nc=sn&loc=2

QUESTION 484
During a security audit of a Service team's application, a Solutions Architect discovers that a username and password for an Amazon RDS database and a set of AWS IAM user credentials can be viewed in the AWS Lambda function code. The Lambda function uses the username and password to run queries on the database, and it uses the IAM credentials to call AWS services in a separate management account.
The Solutions Architect is concerned that the credentials could grant inappropriate access to anyone who can view the Lambda code. The management account and the Service team's account are in separate AWS Organizations organizational units (OUs).
Which combination of changes should the Solutions Architect make to improve the solution's security? (Choose two.)

A. Configure Lambda to assume a role in the management account with appropriate access to AWS.
B. Configure Lambda to use the stored database credentials in AWS Secrets Manager and enable automatic rotation.
C. Create a Lambda function to rotate the credentials every hour by deploying a new Lambda version with the updated credentials.
D. Use an SCP on the management account's OU to prevent IAM users from accessing resources in the Service team's account.
E. Enable AWS Shield Advanced on the management account to shield sensitive resources from unauthorized IAM access.

Correct Answer: BD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 485
In Amazon Cognito, your mobile app authenticates with the Identity Provider (IdP) using the provider's SDK. Once the end user is authenticated with the IdP, the OAuth or OpenID Connect token returned from the IdP is passed by your app to Amazon Cognito, which returns a new _____ for the user and a set of temporary, limited-privilege AWS credentials.

A. Cognito Key Pair

B. Cognito API
C. Cognito ID
D. Cognito SDK

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Your mobile app authenticates with the identity provider (IdP) using the provider's SDK. Once the end user is authenticated with the IdP, the OAuth or OpenID Connect token returned from the IdP is passed by your app to Amazon Cognito, which returns a new Cognito ID for the user and a set of temporary, limited-privilege AWS credentials.
Reference: http://aws.amazon.com/cognito/faqs/

QUESTION 486
A large global company wants to migrate a stateless mission-critical application to AWS. The application is based on IBM WebSphere (application and integration middleware), IBM MQ (messaging middleware), and IBM DB2 (database software) on a z/OS operating system.
How should the Solutions Architect migrate the application to AWS?

A. Re-host WebSphere-based applications on Amazon EC2 behind a load balancer with Auto Scaling. Re-platform the IBM MQ to an Amazon EC2-based MQ. Re-platform the z/OS-based DB2 to Amazon RDS DB2.
B. Re-host WebSphere-based applications on Amazon EC2 behind a load balancer with Auto Scaling. Re-platform the IBM MQ to an Amazon MQ. Re-platform z/OS-based DB2 to Amazon EC2-based DB2.
C. Orchestrate and deploy the application by using AWS Elastic Beanstalk. Re-platform the IBM MQ to Amazon SQS. Replatform z/OS-based DB2 to Amazon RDS DB2.
D. Use the AWS Server Migration Service to migrate the IBM WebSphere and IBM DB2 to an Amazon EC2-based solution. Re-platform the IBM MQ to an Amazon MQ.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://aws.amazon.com/blogs/database/aws-database-migration-service-and-aws-schema-conversion-tool-now-supportibm-db2-as-a-source/ https://aws.amazon.com/quickstart/architecture/ibm-mq/

QUESTION 487

A company uses Amazon S3 to store documents that may only be accessible to an Amazon EC2 instance in a certain virtual private cloud (VPC). The company fears that a malicious insider with access to this instance could also set up an EC2 instance in another VPC to access these documents.
Which of the following solutions will provide the required protection?

A. Use an S3 VPC endpoint and an S3 bucket policy to limit access to this VPC endpoint.
B. Use EC2 instance profiles and an S3 bucket policy to limit access to the role attached to the instance profile.
C. Use S3 client-side encryption and store the key in the instance metadata.
D. Use S3 server-side encryption and protect the key with an encryption context.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 488
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:

QUESTION 489
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 two.)

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: AB
Section: (none)
Explanation

Explanation/Reference:
Explanation:
As per the document defined, new IPs should be reserved not the same ones Elastic IP Addresses are static IP addresses designed for dynamic cloud computing. Unlike traditional static IP addresses, however, Elastic IP addresses enable you to mask instance or Availability Zone failures by programmatically remapping your public IP addresses to instances in your account in a particular region. For DR, you can also pre-allocate some IP addresses for the most critical systems so that their IP addresses are already known before disaster strikes. This can simplify the execution of the DR plan. Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/resources.html

QUESTION 490
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 two.)

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: BE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 491
A large company is running a popular web application. The application runs on several Amazon EC2 Linux instances in an Auto Scaling group in a private subnet. An Application Load Balancer is targeting the instances in the Auto Scaling group in the private subnet. AWS Systems Manager Session Manager is configured, and AWS Systems Manager Agent is running on all the EC2 instances.
The company recently released a new version of the application. Some EC2 instances are now being marked as unhealthy and are being terminated. As a result,

the application is running at reduced capacity. A solutions architect tries to determine the root cause by analyzing Amazon CloudWatch logs that are collected from the application, but the logs are inconclusive.
How should the solutions architect gain access to an EC2 instance to troubleshoot the issue?

A. Suspend the Auto Scaling group's HealthCheck scaling process. Use Session Manager to log in to an instance that is marked as unhealthy.
B. Enable EC2 instance termination protection. Use Session Manager to log in to an instance that is marked as unhealthy.
C. Set the termination policy to OldestInstance on the Auto Scaling group. Use Session Manager to log in to an instance that is marked an unhealthy.
D. Suspend the Auto Scaling group's Terminate process. Use Session Manager to log in to an instance that is marked as unhealthy.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 492
A company has developed a web application that runs on Amazon EC2 instances in one AWS Region. The company has taken on new business in other countries and must deploy its application into other regions to meet low-latency requirements for its users. The regions can be segregated, and an application running in one region does not need to communicate with instances in other regions.
How should the company's Solutions Architect automate the deployment of the application so that it can be MOST efficiently deployed into multiple regions?

A. Write a bash script that uses the AWS CLI to query the current state in one region and output a JSON representation. Pass the JSON representation to the AWS CLI, specifying the --region parameter to deploy the application to other regions.
B. Write a bash script that uses the AWS CLI to query the current state in one region and output an AWS CloudFormation template. Create a CloudFormation stack from the template by using the AWS CLI, specifying the --region parameter to deploy the application to other regions.
C. Write a CloudFormation template describing the application's infrastructure in the resources section. Create a CloudFormation stack from the template by using the AWS CLI, specify multiple regions using the --regions parameter to deploy the application.
D. Write a CloudFormation template describing the application's infrastructure in the Resources section. Use a CloudFormation stack set from an administrator account to launch stack instances that deploy the application to other regions.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A stack set let you create stacks in AWS accounts across regions by using a single AWS CloudFormation template. All the resources included in each stack are defined by the stack set's AWS CloudFormation template. As you create the stack set, you specify the template to use, as well as any parameters and

capabilities that template requires.
Reference: https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-concepts.html
https://sanderknape.com/2017/07/cloudformation-stacksets-automated-cross-account-region-deployments/

QUESTION 493
Which of the following is NOT true of the DynamoDB Console?

A. It allows you to add local secondary indexes to existing tables.
B. It allows you to query a table.
C. It allows you to set up alarms to monitor your table's capacity usage.
D. It allows you to view items stored in a tables, add, update, and delete items.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The DynamoDB Console lets you do the following: Create, update, and delete tables. The throughput calculator provides you with estimates of how many capacity units you will need to request based on the usage information you provide. View items stored in a tables, add, update, and delete items. Query a table. Set up alarms to monitor your table's capacity usage. View your table's top monitoring metrics on real-time graphs from CloudWatch. View alarms configured for each table and create custom alarms.html.

QUESTION 494
A security engineer determined that an existing application retrieves credentials to an Amazon RDS for MySQL database from an encrypted file in Amazon S3. For the next version of the application, the security engineer wants to implement the following application design changes to improve security:
The database must use strong, randomly generated passwords stored in a secure AWS managed service.
The application resources must be deployed through AWS CloudFormation. The application must rotate credentials for the database every 90 days.
A solutions architect will generate a CloudFormation template to deploy the application.
Which resources specified in the CloudFormation template will meet the security engineer's requirements with the LEAST amount of operational overhead?

A. Generate the database password as a secret resource using AWS Secrets Manager. Create an AWS Lambda function resource to rotate the database password. Specify a Secrets Manager RotationSchedule resource to rotate the database password every 90 days.
B. Generate the database password as a SecureString parameter type using AWS Systems Manager Parameter Store. Create an AWS Lambda function resource to rotate the database password. Specify a Parameter Store RotationSchedule resource to rotate the database password every 90 days.
C. Generate the database password as a secret resource using AWS Secrets Manager. Create an AWS Lambda function resource to rotate the database password. Create an Amazon EventBridge scheduled rule resource to trigger the Lambda function password rotation every 90 days.
D. Generate the database password as a SecureString parameter type using AWS Systems Manager Parameter Store. Specify an AWS AppSync DataSource resource to automatically rotate the database password every 90 days.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 495
An organization is undergoing a security audit. The auditor wants to view the AWS VPC configurations as the organization has hosted all the applications in the AWS VPC. The auditor is from a remote place and wants to have access to AWS to view all the VPC records.
How can the organization meet the expectations of the auditor without compromising on the security of their AWS infrastructure?

A. The organization should not accept the request as sharing the credentials means compromising on security.
B. Create an IAM role which will have read only access to all EC2 services including VPC and assign that role to the auditor.
C. Create an IAM user who will have read only access to the AWS VPC and share those credentials with the auditor.
D. The organization should create an IAM user with VPC full access but set a condition that will not allow to modify anything if the request is from any IP other than the organization's data center.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. The VPC also works with IAM and the organization can create IAM users who have access to various VPC services. If an auditor wants to have access to the AWS VPC to verify the rules, the organization should be careful before sharing any data which can allow making updates to the AWS infrastructure. In this scenario it is recommended that the organization creates an IAM user who will have read only access to the VPC. Share the above mentioned credentials with the auditor as it cannot harm the organization. The sample policy is given below:
{
"Effect":"Allow", "Action": ["ec2:DescribeVpcs", "ec2:DescribeSubnets", "ec2: DescribeInternetGateways", "ec2:DescribeCustomerGateways", "ec2:DescribeVpnGateways", "ec2:DescribeVpnConnections", "ec2:DescribeRouteTables", "ec2:DescribeAddresses", "ec2:DescribeSecurityGroups", "ec2:DescribeNetworkAcls", "ec2:DescribeDhcpOptions", "ec2:DescribeTags", "ec2:DescribeInstances"], "Resource":"*"
}
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_IAM.html

QUESTION 496
An organization hosts an app on EC2 instances which multiple developers need access to in order to perform updates. The organization plans to implement some security best practices related to instance access.

Which one of the following recommendations will not help improve its security in this way?

A. Disable the password based login for all the users. All the users should use their own keys to connect with the instance securely.
B. Create an IAM policy allowing only IAM users to connect to the EC2 instances with their own SSH key.
C. Create a procedure to revoke the access rights of the individual user when they are not required to connect to EC2 instance anymore for the purpose of application configuration.
D. Apply the latest patch of OS and always keep it updated.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Since AWS is a public cloud any application hosted on EC2 is prone to hacker attacks. It becomes extremely important for a user to setup a proper security mechanism on the EC2 instances. A few of the security measures are listed below:
Always keep the OS updated with the latest patch
Always create separate users with in OS if they need to connect with the EC2 instances, create their keys and disable their password Create a procedure using which the admin can revoke the access of the user when the business work on the EC2 instance is completed. . Lock down unnecessary ports Audit any proprietary applications that the user may be running on the EC2 instance. Provide temporary escalated privileges, such as sudo for users who need to perform occasional privileged tasks IAM is useful when users are required to work with AWS resources and actions, such as launching an instance. It is not useful in this case because it does not manage who can connect via RDP or SSH with an instance.
Reference: http://aws.amazon.com/articles/1233/

QUESTION 497

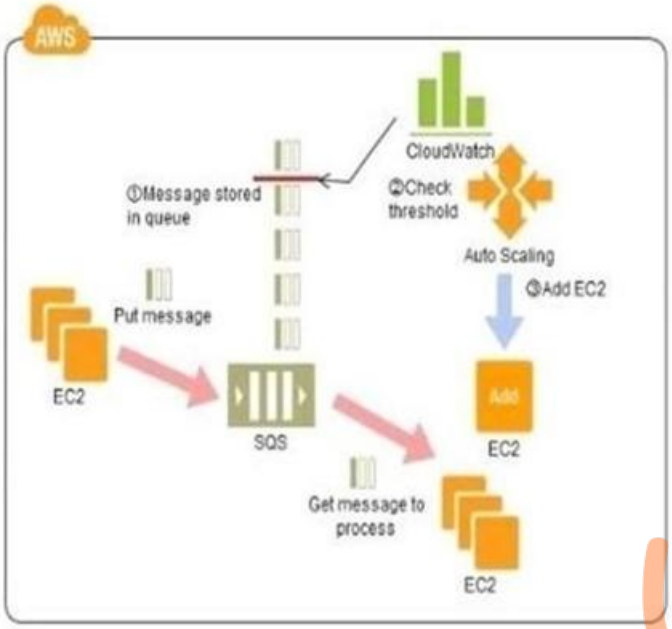

Refer to the architecture diagram above of a batch processing solution using Simple Queue Service (SQS) to set up a message queue between EC2 instances which are used as batch processors Cloud Watch monitors the number of Job requests (queued messages) and an Auto Scaling group adds or deletes batch 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 reprioritized.
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?
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

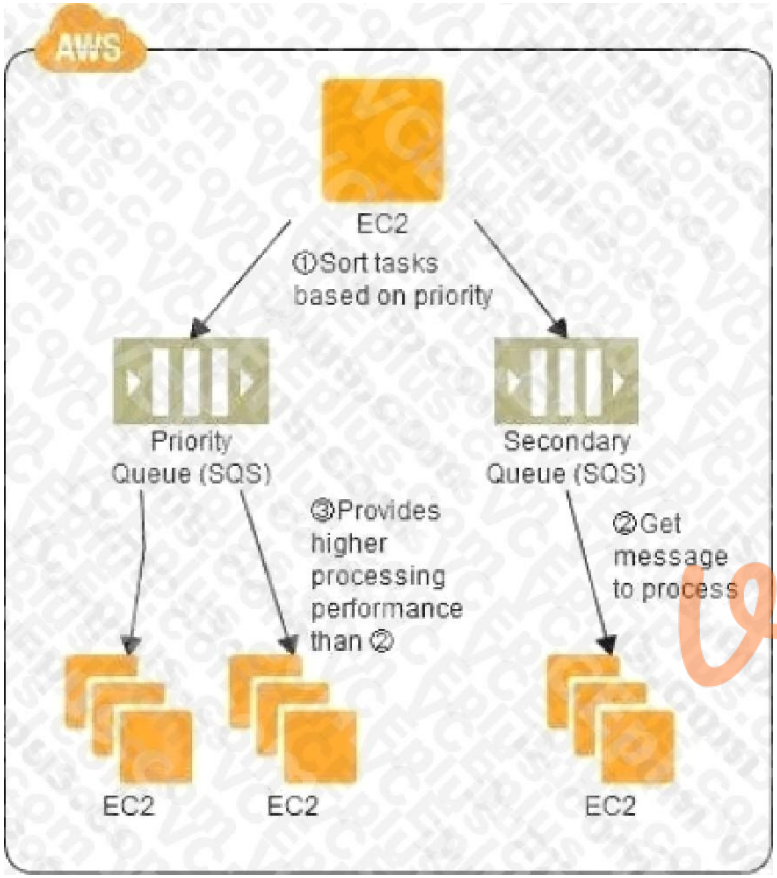

Benefits

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 498

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 in different Availability Zones for greater levels of data durability. For more information on Multi-AZ deployments, see High Availability (Multi-AZ).

QUESTION 499
You have written a CloudFormation template that creates 1 Elastic Load Balancer fronting 2 EC2 Instances.
Which section of the template should you edit so that the DNS of the load balancer is returned upon creation of the stack?

A. Parameters
B. Outputs
C. Mappings
D. Resources

Correct Answer: B

Section: (none)
Explanation

Explanation/Reference:
Explanation:
You can use AWS CloudFormation's sample templates or create your own templates to describe the AWS resources, and any associated dependencies or runtime parameters, required to run your application.
In the following example, the output named BackupLoadBalancerDNSName returns the DNS name for the resource with the logical ID BackupLoadBalancer only when the CreateProdResources condition is true. (The second output shows how to specify multiple outputs.) "Outputs" : { "BackupLoadBalancerDNSName" : {
"Description": "The DNSName of the backup load balancer", "Value" : { "Fn::GetAtt" : ["BackupLoadBalancer", "DNSName"
]}, "Condition" : "CreateProdResources"
},
"InstanceID" : {
"Description": "The Instance ID", "Value" : { "Ref" : "EC2Instance" }
}
}
Reference: http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/outputs-section-structure.html

QUESTION 500
A finance company is running its business-critical application on current-generation Linux EC2 instances. The application includes a self-managed MySQL database performing heavy I/O operations. The application is working fine to handle a moderate amount of traffic during the month. However, it slows down during the final three days of each month due to month-end reporting, even though the company is using Elastic Load Balancers and Auto Scaling within its infrastructure to meet the increased demand.
Which of the following actions would allow the database to handle the month-end load with the LEAST impact on performance?

A. Pre-warming Elastic Load Balancers, using a bigger instance type, changing all Amazon EBS volumes to GP2 volumes.
B. Performing a one-time migration of the database cluster to Amazon RDS, and creating several additional read replicas to handle the load during end of month.
C. Using Amazon CloudWatch with AWS Lambda to change the type, size, or IOPS of Amazon EBS volumes in the cluster based on a specific CloudWatch metric.
D. Replacing all existing Amazon EBS volumes with new PIOPS volumes that have the maximum available storage size and I/O per second by taking snapshots before the end of the month and reverting back afterwards.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 501
Does an AWS Direct Connect location provide access to Amazon Web Services in the region it is associated with as well as access to other US regions?

A. No, it provides access only to the region it is associated with.
B. No, it provides access only to the US regions other than the region it is associated with.
C. Yes, it provides access.
D. Yes, it provides access but only when there's just one Availability Zone in the region.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
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).
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.html

QUESTION 502
A financial services company logs personally identifiable information to its application logs stored in Amazon S3. Due to regulatory compliance requirements, the log files must be encrypted at rest. The security team has mandated that the company's on-premises hardware security modules (HSMs) be used to generate the CMK material.
Which steps should the solutions architect take to meet these requirements?

A. Create an AWS CloudHSM cluster. Create a new CMK in AWS KMS using AWS_CloudHSM as the source for the key material and an origin of AWS_CLOUDHSM. Enable automatic key rotation on the CMK with a duration of 1 year. Configure a bucket policy on the logging bucket that disallows uploads of unencrypted data and requires that the encryption source be AWS KMS.
B. Provision an AWS Direct Connect connection, ensuring there is no overlap of the RFC 1918 address space between onpremises hardware and the VPCs. Configure an AWS bucket policy on the logging bucket that requires all objects to be encrypted. Configure the logging application to query the on-premises HSMs from the AWS environment for the encryption key material, and create a unique CMK for each logging event.
C. Create a CMK in AWS KMS with no key material and an origin of EXTERNAL. Import the key material generated from the on-premises HSMs into the CMK using the public key and import token provided by AWS. Configure a bucket policy on the logging bucket that disallows uploads of non-encrypted data and requires that the encryption source be AWS KMS.
D. Create a new CMK in AWS KMS with AWS-provided key material and an origin of AWS_KMS. Disable this CMK, and overwrite the key material with the key material from the on-premises HSM using the public key and import token provided by AWS. Re-enable the CMK. Enable automatic key rotation on the CMK with a duration of 1 year. Configure a bucket policy on the logging bucket that disallows uploads of non-encrypted data and requires that the encryption source be AWS KMS.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 503
A scientific company needs to process text and image data from an Amazon S3 bucket. The data is collected from several radar stations during a live, time-critical phase of a deep space mission. The radar stations upload the data to the source S3 bucket. The data is prefixed by radar station identification number. The company created a destination S3 bucket in a second account. Data must be copied from the source S3 bucket to the destination S3 bucket to meet a compliance objective. The replication occurs through the use of an S3 replication rule to cover all objects in the source S3 bucket.
One specific radar station is identified as having the most accurate data. Data replication at this radar station must be monitored for completion within 30 minutes after the radar station uploads the objects to the source S3 bucket.
What should a solutions architect do to meet these requirements?

A. Set up an AWS DataSync agent to replicate the prefixed data from the source S3 bucket to the destination S3 bucket.
Select to use all available bandwidth on the task, and monitor the task to ensure that it is in the TRANSFERRING status.
Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert if this status changes.
B. In the second account, create another S3 bucket to receive data from the radar station with the most accurate data. Set up a new replication rule for this new S3 bucket to separate the replication from the other radar stations. Monitor the maximum replication time to the destination. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert when the time exceeds the desired threshold.
C. Enable Amazon S3 Transfer Acceleration on the source S3 bucket, and configure the radar station with the most accurate data to use the new endpoint. Monitor the S3 destination bucket's TotalRequestLatency metric. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert if this status changes.
D. Create a new S3 replication rule on the source S3 bucket that filters for the keys that use the prefix of the radar station with the most accurate data. Enable S3 Replication Time Control (S3 RTC). Monitor the maximum replication time to the destination. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert when the time exceeds the desired threshold.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 504
A financial company is building a system to generate monthly, immutable bank account statements for its users. Statements are stored in Amazon S3. Users should have immediate access to their monthly statements for up to 2 years. Some users access their statements frequently, whereas others rarely access their statements. The company's security and compliance policy requires that the statements be retained for at least 7 years.

What is the MOST cost-effective solution to meet the company's needs?

A. Create an S3 bucket with Object Lock disabled. Store statements in S3 Standard. Define an S3 Lifecycle policy to transition the data to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days. Define another S3 Lifecycle policy to move the data to S3 Glacier Deep Archive after 2 years. Attach an S3 Glacier Vault Lock policy with deny delete permissions for archives less than 7 years old.
B. Create an S3 bucket with versioning enabled. Store statements in S3 Intelligent-Tiering. Use same-Region replication to replicate objects to a backup S3 bucket. Define an S3 Lifecycle policy for the backup S3 bucket to move the data to S3 Glacier. Attach an S3 Glacier Vault Lock policy with deny delete permissions for archives less than 7 years old.
C. Create an S3 bucket with Object Lock enabled. Store statements in S3 Intelligent-Tiering. Enable compliance mode with a default retention period of 2 years. Define an S3 Lifecycle policy to move the data to S3 Glacier after 2 years. Attach an S3 Glacier Vault Lock policy with deny delete permissions for archives less than 7 years old.
D. Create an S3 bucket with versioning disabled. Store statements in S3 One Zone-Infrequent Access (S3 One Zone-IA). Define an S3 Lifecycle policy to move the data to S3 Glacier Deep Archive after 2 years. Attach an S3 Glacier Vault Lock policy with deny delete permissions for archives less than 7 years old.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 505
In order for a table write to succeed, the provisioned throughput settings for the table and global secondary indexes, in DynamoDB, must have_____; otherwise, the write to the table will be throttled.

A. enough write capacity to accommodate the write
B. no additional write cost for the index
C. 100 bytes of overhead per index item
D. the size less than or equal to 1 KB

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In order for a table write to succeed in DynamoDB, the provisioned throughput settings for the table and global secondary indexes must have enough write capacity to accommodate the write; otherwise, the write will be throttled.

Reference: http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GSI.html

QUESTION 506
You have set up Auto Scaling to automatically scale in. Consequently, you must decide which instances Auto Scaling should end first.
What should you use to configure this?

A. An Elastic Load Balancer
B. A termination policy
C. An IAM role
D. Another scaling group

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
If you instruct Auto Scaling to automatically scale in, you must decide which instances Auto Scaling should terminate first.
This can be configured through the use of a termination policy.
Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AutoScalingBehavior.InstanceTermination.html

QUESTION 507
A company is having issues with a newly deployed serverless infrastructure that uses Amazon API Gateway, Amazon Lambda, and Amazon DynamoDB.
In a steady state, the application performs as expected. However, during peak load, tens of thousands of simultaneous invocations are needed and user requests fail multiple times before succeeding. The company has checked the logs for each component, focusing specifically on Amazon CloudWatch Logs for Lambda.
There are no errors logged by the services or applications.
What might cause this problem?

A. Lambda has very low memory assigned, which causes the function to fail at peak load.
B. Lambda is in a subnet that uses a NAT gateway to reach out of the internet, and the function instance does not have sufficient Amazon EC2 resources in the VPC to scale with the load.
C. The throttle limit set on API Gateway is very low. During peak load, the additional requests are not making their way through to Lambda.
D. DynamoDB is set up in an auto scaling mode. During peak load, DynamoDB adjusts capacity and throughput behind the scenes, which is causing the temporary downtime. Once the scaling completes, the retries go through successfully.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-request-throttling.html

QUESTION 508
Will you be able to access EC2 snapshots using the regular Amazon S3 APIs?

A. Yes, you will be able to access using S3 APIs if you have chosen the snapshot to be stored in S3.
B. No, snapshots are only available through the Amazon EBS APIs.
C. Yes, you will be able to access them using S3 APIs as all snapshots are stored in S3.
D. No, snapshots are only available through the Amazon EC2 APIs.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
No, snapshots are only available through the Amazon EC2 APIs.
Reference: https://aws.amazon.com/ec2/faqs/

QUESTION 509
Does Amazon RDS API provide actions to modify DB instances inside a VPC and associate them with DB Security Groups?

A. Yes, Amazon does this but only for MySQL RDS.
B. Yes
C. No
D. Yes, Amazon does this but only for Oracle RDS.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
You can use the action Modify DB Instance, available in the Amazon RDS API, to pass values for the parameters DB Instance Identifier and DB Security Groups specifying the instance ID and the DB Security Groups you want your instance to be part of.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/APIReference/API_ModifyDBInstance.html

QUESTION 510
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 511
A user is sending a custom metric to CloudWatch. If the call to the CloudWatch APIs has different dimensions, but the same metric name, how will CloudWatch treat all the requests?

A. It will reject the request as there cannot be a separate dimension for a single metric.
B. It will group all the calls into a single call.
C. It will treat each unique combination of dimensions as a separate metric.
D. It will overwrite the previous dimension data with the new dimension data.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

A dimension is a key-value pair used to uniquely identify a metric. CloudWatch treats each unique combination of dimensions as a separate metric. Thus, if the user is making 4 calls with the same metric name but a separate dimension, it will create 4 separate metrics.
Reference: http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/cloudwatch_concepts.html

QUESTION 512
A user is planning to host a Highly Available system on the AWS VPC. Which of the below mentioned statements is helpful in this scenario?

A. Create VPC subnets in two separate availability zones and launch instances in different subnets.
B. Create VPC with only one public subnet and launch instances in different AZs using that subnet.
C. Create two VPCs in two separate zones and setup failover with ELB such that if one VPC fails it will divert traffic to another VPC.
D. Create VPC with only one private subnet and launch instances in different AZs using that subnet.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. The VPC is always specific to a region. The user can create a VPC which can span multiple Availability Zones by adding one or more subnets in each Availability Zone. Each subnet must reside entirely within one Availability Zone and cannot span across zones.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html#VPCSubnet

QUESTION 513
A company's main intranet page has experienced degraded response times as its user base has increased although there are no reports of users seeing error pages. The application uses Amazon DynamoDB in read-only mode.
Amazon DynamoDB latency metrics for successful requests have been in a steady state even during times when users have reported degradation. The Development team has correlated the issue to ProvisionedThrough put Exceeded exceptions in the application logs when doing Scan and read operations The team also identified an access pattern of steady spikes of read activity on a distributed set of individual data items.
The Chief Technology Officer wants to improve the user experience.
Which solutions will meet these requirements with the LEAST amount of changes to the application? (Choose two.)

A. Change the data model of the DynamoDB tables to ensure that all Scan and read operations meet DynamoDB best practices of uniform data access, reaching the full request throughput provisioned for the DynamoDB tables.
B. Enable DynamoDB Auto Scaling to manage the throughput capacity as table traffic increases. Set the upper and lower limits to control costs and set a target utilization given the peak usage and how quickly the traffic changes.
C. Provision Amazon ElastiCache for Redis with cluster mode enabled. The cluster should be provisioned with enough shards to spread the application load and provision at least one read replica node for each shard.
D. Implement the DynamoDB Accelerator (DAX) client and provision a DAX cluster with the appropriate node types to sustain the application load. Tune the item and query cache configuration for an optimal user experience.

E. Remove error retries and exponential backoffs in the application code to handle throttling errors.

Correct Answer: AE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 514
A company has several Amazon EC2 instances to both public and private subnets within a VPC that is not connected to the corporate network. A security group associated with the EC2 instances allows the company to use the Windows remote desktop protocol (RDP) over the internet to access the instances. The security team has noticed connection attempts from unknown sources. The company wants to implement a more secure solution to access the EC2 instances.
Which strategy should a solutions architect implement?

A. Deploy a Linux bastion host on the corporate network that has access to all instances in the VPC.
B. Deploy AWS Systems Manager Agent on the EC2 instances. Access the EC2 instances using Session Manager restricting access to users with permission.
C. Deploy a Linux bastion host with an Elastic IP address in the public subnet. Allow access to the bastion host from 0.0.0.0/0.
D. Establish a Site-to-Site VPN connecting the corporate network to the VPC. Update the security groups to allow access from the corporate network only.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 515
A company needs to establish a connection from its on-premises data center to AWS. The company needs to connect all of its VPCs that are located in different AWS Regions with transitive routing capabilities between VPC networks. The company also must reduce network outbound traffic costs, increase bandwidth throughput, and provide a consistent network experience for end users.
Which solution will meet these requirements?

A. Create an AWS Site-to-Site VPN connection between the on-premises data center and a new central VPC. Create VPC peering connections that initiate from the central VPC to all other VPCs.
B. Create an AWS Direct Connect connection between the on-premises data center and AWS. Provision a transit VIF, and connect it to a Direct Connect gateway. Connect the Direct Connect gateway to all the other VPCs by using a transit gateway in each Region.
C. Create an AWS Site-to-Site VPN connection between the on-premises data center and a new central VPUse a transit gateway with dynamic routing. Connect the transit gateway to all other VPCs.

D. Create an AWS Direct Connect connection between the on-premises data center and AWS. Establish an AWS Site-to- Site VPN connection between all VPCs in each Region. Create VPC peering connections that initiate from the central VPC to all other VPCs.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/vpc/faqs/

QUESTION 516
You create a VPN connection, and your VPN device supports Border Gateway Protocol (BGP).
Which of the following should be specified to configure the VPN connection?

A. Classless routing
B. Classfull routing
C. Dynamic routing
D. Static routing

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
If you create a VPN connection, you must specify the type of routing that you plan to use, which will depend upon on the make and model of your VPN devices. If your VPN device supports Border Gateway Protocol (BGP), you need to specify dynamic routing when you configure your VPN connection. If your device does not support BGP, you should specify static routing.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_VPN.html

QUESTION 517
A solutions architect has implemented a SAML 2.0 federated identity solution with their company's on-premises identity provider (IdP) to authenticate users' access to the AWS environment. When the solutions architect tests authentication through the federated identity web portal, access to the AWS environment is granted. However, when test users attempt to authenticate through the federated identity web portal, they are not able to access the AWS environment.
Which items should the solutions architect check to ensure identity federation is properly configured? (Choose three.)

A. The IAM user's permissions policy has allowed the use of SAML federation for that user.
B. The IAM roles created for the federated users' or federated groups' trust policy have set the SAML provider as the principal.

C. Test users are not in the AWSFederatedUsers group in the company's IdR.
D. The web portal calls the AWS STS AssumeRoleWithSAML API with the ARN of the SAML provider, the ARN of the IAM role, and the SAML assertion from IdR.
E. The on-premises IdP's DNS hostname is reachable from the AWS environment VPCs.
F. The company's IdP defines SAML assertions that properly map users or groups in the company to IAM roles with appropriate permissions.

Correct Answer: DEF
Section: (none)
Explanation

Explanation/Reference:

QUESTION 518
An organization is setting up a multi-site solution where the application runs on premise as well as on AWS to achieve the minimum recovery time objective (RTO).
Which of the below mentioned configurations will not meet the requirements of the multi-site solution scenario?

A. Configure data replication based on RTO.
B. Keep an application running on premise as well as in AWS with full capacity.
C. Setup a single DB instance which will be accessed by both sites.
D. Setup a weighted DNS service like Route 53 to route traffic across sites.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 519
An organization is trying to setup a VPC with Auto Scaling. Which configuration steps below is not required to setup AWS VPC with Auto Scaling?

A. Configure the Auto Scaling group with the VPC ID in which instances will be launched.
B. Configure the Auto Scaling Launch configuration with multiple subnets of the VPC to enable the Multi AZ feature.
C. Configure the Auto Scaling Launch configuration which does not allow assigning a public IP to instances.
D. Configure the Auto Scaling Launch configuration with the VPC security group.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Within this virtual private cloud, the user can launch AWS resources, such as an Auto Scaling group. Before creating the Auto Scaling group it is recommended that the user creates the Launch configuration. Since it is a VPC, it is recommended to select the parameter which does not allow assigning a public IP to the instances.
The user should also set the VPC security group with the Launch configuration and select the subnets where the instances will be launched in the AutoScaling group. The HA will be provided as the subnets may be a part of separate AZs.
Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/autoscalingsubnets.html

QUESTION 520
A travel company built a web application that uses Amazon Simple Email Service (Amazon SES) to send email notifications to users. The company needs to enable logging to help troubleshoot email delivery issues. The company also needs the ability to do searches that are based on recipient, subject, and time sent.
Which combination of steps should a solutions architect take to meet these requirements? (Choose two.)

A. Create an Amazon SES configuration set with Amazon Kinesis Data Firehose as the destination. Choose to send logs to an Amazon S3 bucket.
B. Enable AWS CloudTrail logging. Specify an Amazon S3 bucket as the destination for the logs.
C. Use Amazon Athena to query the logs in the Amazon S3 bucket for recipient, subject, and time sent.
D. Create an Amazon CloudWatch log group. Configure Amazon SES to send logs to the log group.
E. Use Amazon Athena to query the logs in Amazon CloudWatch for recipient, subject, and time sent.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference https://docs.aws.amazon.com/ses/latest/DeveloperGuide/ses-dg.pdf

QUESTION 521
A company used Amazon EC2 instances to deploy a web fleet to host a blog site. The EC2 instances are behind an Application Load Balancer (ALB) and are configured in an Auto Scaling group. The web application stores all blog content on an Amazon EFS volume.
The company recently added a feature for bloggers to add video to their posts, attracting 10 times the previous user traffic.
At peak times of day, users report buffering and timeout issues while attempting to reach the site or watch videos.
Which is the MOST cost-efficient and scalable deployment that will resolve the issues for users?

A. Reconfigure Amazon EFS to enable maximum I/O.
B. Update the blog site to use instance store volumes for storage. Copy the site contents to the volumes at launch and to Amazon S3 at shutdown.
C. Configure an Amazon CloudFront distribution. Point the distribution to an S3 bucket, and migrate the videos from EFS to Amazon S3.
D. Set up an Amazon CloudFront distribution for all site contents, and point the distribution at the ALB.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-https-connection-fails/

QUESTION 522
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 two.)

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: CD
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://d0.awsstatic.com/whitepapers/performance-at-scale-with-amazon-elasticache.pdf

QUESTION 523

Which of the following does Amazon DynamoDB perform?

A. Atomic increment or decrement on scalar values
B. Neither increment nor decrement operations
C. Only increment on vector values
D. Only atomic decrement operations

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Amazon DynamoDB allows atomic increment and decrement operations on scalar values.
Reference: http://aws.amazon.com/dynamodb/faqs/

QUESTION 524
What happens when Dedicated instances are launched into a VPC?

A. If you launch an instance into a VPC that has an instance tenancy of dedicated, you must manually create a Dedicated instance.
B. If you launch an instance into a VPC that has an instance tenancy of dedicated, your instance is created as a Dedicated instance, only based on the tenancy of the instance.
C. If you launch an instance into a VPC that has an instance tenancy of dedicated, your instance is automatically a Dedicated instance, regardless of the tenancy of the instance.
D. None of these are true.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
If you launch an instance into a VPC that has an instance tenancy of dedicated, your instance is automatically a Dedicated instance, regardless of the tenancy of the instance.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/dedicated-instance.html

QUESTION 525
A company is deploying a third-party firewall appliance solution from AWS Marketplace to monitor and protect traffic that leaves the company's AWS environments. The company wants to deploy this appliance into a shared services VPC and route all outbound internet-bound traffic through the appliances.

A solutions architect needs to recommend a deployment method that prioritizes reliability and minimizes failover time between firewall appliances within a single AWS Region. The company has set up routing from the shared services VPC to other VPCs.
Which steps should the solutions architect recommend to meet these requirements? (Choose three.)

A. Deploy two firewall appliances into the shared services VPC, each in a separate Availability Zone.
B. Create a new Network Load Balancer in the shared services VPC. Create a new target group, and attach it to the new Network Load Balancer. Add each of the firewall appliance instances to the target group.
C. Create a new Gateway Load Balancer in the shared services VPCreate a new target group, and attach it to the new Gateway Load Balancer. Add each of the firewall appliance instances to the target group.
D. Create a VPC interface endpoint. Add a route to the route table in the shared services VPC. Designate the new endpoint as the next hop for traffic that enters the shared services VPC from other VPCs.
E. Deploy two firewall appliances into the shared services VPC, each in the same Availability Zone.
F. Create a VPC Gateway Load Balancer endpoint. Add a route to the route table in the shared services VPC. Designate the new endpoint as the next hop for traffic that enters the shared services VPC from other VPCs.

Correct Answer: BEF
Section: (none)
Explanation

Explanation/Reference:

QUESTION 526
A large global financial services company has multiple business units. The company wants to allow Developers to try new services, but there are multiple compliance requirements for different workloads. The Security team is concerned about the access strategy for on-premises and AWS implementations. They would like to enforce governance for AWS services used by business teams for regulatory workloads, including Payment Card Industry (PCI) requirements.
Which solution will address the Security team's concerns and allow the Developers to try new services?

A. Implement a strong identity and access management model that includes users, groups, and roles in various AWS accounts. Ensure that centralized AWS CloudTrail logging is enabled to detect anomalies. Build automation with AWS Lambda to tear down unapproved AWS resources for governance.
B. Build a multi-account strategy based on business units, environments, and specific regulatory requirements. Implement SAML-based federation across all AWS accounts with an on-premises identity store. Use AWS Organizations and build organizational units (OUs) structure based on regulations and service governance. Implement service control policies across OUs.
C. Implement a multi-account strategy based on business units, environments, and specific regulatory requirements. Ensure that only PCI-compliant services are approved for use in the accounts. Build IAM policies to give access to only PCIcompliant services for governance.
D. Build one AWS account for the company for strong security controls. Ensure that all the service limits are raised to meet company scalability requirements. Implement SAML federation with an onpremises identity store, and ensure that only approved services are used in the account.

Correct Answer: B

Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_introduction.html

QUESTION 527
When I/O performance is more important than fault tolerance, which of the following configurations should be used?

A. SPAN 10
B. RAID 1
C. RAID 0
D. NFS 1

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
When I/O performance is more important than fault tolerance, the RAID 0 configuration must be used; for example, as in a heavily used database (where data replication is already set up separately).
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/raid-config.html

QUESTION 528
A company asks a solution architect to optimize the cost of a solution. The solution handles requests from multiple customers. The solution includes a multi-tier architecture that uses Amazon API Gateway, AWS Lambda, AWS Fargate, Amazon Simple Queue Service (Amazon SQS), and Amazon EC2.
In the current setup, requests go through API Gateway to Lambda and either start a container in Fargate or push a message to an SQS queue. An EC2 Fleet provides EC2 instances that serve as workers for the SQS queue. The EC2 Fleet scales based on the number of items in the SQS queue.
Which combination of steps should the solutions architect recommend to reduce cost the MOST? (Choose three.)

A. Determine the minimum number of EC2 instances that are needed during a day. Reserve this number of instances in a 3- year plan with payment all upfront.
B. Examine the last 6 months of compute utilization across the services. Use this information to determine the needed compute for the solution. Commit to a Savings Plan for this amount.
C. Determine the average number of EC2 instances that are needed during a day. Reserve this number of instances in a 3- year plan with payment all upfront.
D. Remove the SQS queue from the solution and from the solution infrastructure.
E. Change the solution so that it runs as a container instead of on EC2 instances. Configure Lambda to start up the solution in Fargate by using environment

variables to give the solution the message.

F. Change the Lambda function so that it posts the message directly to the EC2 instances through an Application Load Balancer.

Correct Answer: CDE
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/ec2/pricing/reserved-instances/

QUESTION 529
You're trying to delete an SSL certificate from the IAM certificate store, and you're getting the message "Certificate: is being used by CloudFront." Which of the following statements is probably the reason why you are getting this error?

A. Before you can delete an SSL certificate you need to set up https on your server.

B. Before you can delete an SSL certificate, you need to set up the appropriate access level in IAM

C. Before you can delete an SSL certificate, you need to either rotate SSL certificates or revert from using a custom SSL certificate to using the default CloudFront certificate.

D. You can't delete SSL certificates. You need to request it from AWS.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
CloudFront is a web service that speeds up distribution of your static and dynamic web content, for example, .html, .css,.php, and image files, to end users. Every CloudFront web distribution must be associated either with the default CloudFront certificate or with a custom SSL certificate. Before you can delete an SSL certificate, you need to either rotate SSL certificates (replace the current custom SSL certificate with another custom SSL certificate) or revert from using a custom SSL certificate to using the default CloudFront certificate.
Reference: http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Troubleshooting.html

QUESTION 530
Which of the following is the final step that should be completed to start using AWS Direct Connect?

A. Creating your Virtual Interface

B. Configuring your router

C. Completing the Cross Connect
D. Verifying your Virtual Interface

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
You can get started using AWS Direct Connect by completing the following steps. Step 1: Sign Up for Amazon Web Services Step 2: Submit AWS Direct Connect Connection Request Step 3: Complete the Cross Connect (optional) Step 4: Configure Redundant Connections with AWS Direct Connect Step 5: Create a Virtual Interface Step 6: Download Router Configuration Step 7: Verify Your Virtual Interface
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/getstarted.html#connected

QUESTION 531
A company is running a large containerized workload in the AWS Cloud. The workload consists of approximately 100 different services. The company uses Amazon Elastic Container Service (Amazon ECS) to orchestrate the workload.
Recently, the company's development team started using AWS Fargate instead of Amazon EC2 instances in the ECS cluster. In the past, the workload has come close to running the maximum number of EC2 instances that are available in the account.
The company is worried that the workload could reach the maximum number of ECS tasks that are allowed. A solutions architect must implement a solution that will notify the development team when Fargate reaches 80% of the maximum number of tasks.
What should the solutions architect do to meet this requirement?

A. Use Amazon CloudWatch to monitor the Sample Count statistic for each service in the ECS cluster. Set an alarm for when the math expression sample count/ SERVICE_QUOTA(service)*100 is greater than 80. Notify the development team by using Amazon Simple Notification Service (Amazon SNS).
B. Use Amazon CloudWatch to monitor service quotas that are published under the AWS/Usage metric namespace. Set an alarm for when the math expression metric/SERVICE_QUOTA(metric)*100 is greater than 80. Notify the development team by using Amazon Simple Notification Service (Amazon SNS).
C. Create an AWS Lambda function to poll detailed metrics form the ECS cluster. When the number of running Fargate tasks is greater than 80, invoke Amazon Simple Email Service (Amazon SES) to notify the development team.
D. Create an AWS Config rule to evaluate whether the Fargate SERVICE_QUOTA is greater than 80. Use Amazon Simple Email Service (Amazon SES) to notify the development team when the AWS Config rule is not compliant.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
To visualize a service quota and optionally set an alarm.
Reference: https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/CloudWatch-Quotas-Visualize-Alarms.html

QUESTION 532
An IAM user is trying to perform an action on an object belonging to some other root account's bucket.
Which of the below mentioned options will AWS S3 not verify?

A. The object owner has provided access to the IAM user
B. Permission provided by the parent of the IAM user on the bucket
C. Permission provided by the bucket owner to the IAM user
D. Permission provided by the parent of the IAM user

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
If the IAM user is trying to perform some action on the object belonging to another AWS user's bucket, S3 will verify whether the owner of the IAM user has given sufficient permission to him. It also verifies the policy for the bucket as well as the policy defined by the object owner.
Reference: http://docs.aws.amazon.com/AmazonS3/latest/dev/access-control-auth-workflow-object-operation.html

QUESTION 533
A company has a Microsoft SQL Server database in its data center and plans to migrate data to Amazon Aurora MySQL.
The company has already used the AWS Schema Conversion Tool to migrate triggers, stored procedures and other schema objects to Aurora MySQL. The database contains 1 TB of data and grows less than 1 MB per day. The company's data center is connected to AWS through a dedicated 1Gbps AWS Direct Connect connection.
The company would like to migrate data to Aurora MySQL and perform reconfigurations with minimal downtime to the applications.
Which solution meets the company's requirements?

A. Shut down applications over the weekend. Create an AWS DMS replication instance and task to migrate existing data from SQL Server to Aurora MySQL. Perform application testing and migrate the data to the new database endpoint.
B. Create an AWS DMS replication instance and task to migrate existing data and ongoing replication from SQL Server to Aurora MySQL. Perform application testing and migrate the data to the new database endpoint.
C. Create a database snapshot of SQL Server on Amazon S3. Restore the database snapshot from Amazon S3 to Aurora MySQL. Create an AWS DMS replication instance and task for ongoing replication from SQL Server to Aurora MySQL.
Perform application testing and migrate the data to the new database endpoint.
D. Create a SQL Server native backup file on Amazon S3. Create an AWS DMS replication instance and task to restore the SQL Server backup file to Aurora MySQL. Create another AWS DMS task for ongoing replication from SQL Server to Aurora MySQL. Perform application testing and migrate the data to the new database endpoint.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 534
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:
Explanation:
Reference: https://aws.amazon.com/developertools/2759763385083070

QUESTION 535
A company is hosting a three-tier web application in an on-premises environment. Due to a recent surge in traffic that resulted in downtime and a significant financial impact, company management has ordered that the application be moved to AWS. The application is written in .NET and has a dependency on a MySQL database. A solutions architect must design a scalable and highly available solution to meet the demand of 200,000 daily users.
Which steps should the solutions architect take to design an appropriate solution?

A. Use AWS Elastic Beanstalk to create a new application with a web server environment and an Amazon RDS MySQL Multi-AZ DB instance. The environment should launch a Network Load Balancer (NLB) in front of an Amazon EC2 Auto Scaling group in multiple Availability Zones. Use an Amazon Route 53 alias record to route traffic from the company's domain to the NLB.
B. Use AWS CloudFormation to launch a stack containing an Application Load Balancer (ALB) in front of an Amazon EC2 Auto Scaling group spanning three Availability Zones. The stack should launch a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a Retain deletion policy. Use an Amazon

Route 53 alias record to route traffic from the company's domain to the ALB.

C. Use AWS Elastic Beanstalk to create an automatically scaling web server environment that spans two separate Regions with an Application Load Balancer (ALB) in each Region. Create a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a cross-Region read replica. Use Amazon Route 53 with a geoproximity routing policy to route traffic between the two Regions.

D. Use AWS CloudFormation to launch a stack containing an Application Load Balancer (ALB) in front of an Amazon ECS cluster of Spot instances spanning three Availability Zones. The stack should launch an Amazon RDS MySQL DB instance with a Snapshot deletion policy. Use an Amazon Route 53 alias record to route traffic from the company's domain to the ALB.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 536
When using the AWS CLI for AWS CloudFormation, which of the following commands returns a description of the specified resource in the specified stack?

A. describe-stack-events
B. describe-stack-resource
C. create-stack-resource
D. describe-stack-returns

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation: awsclicloudformation describe-stack-resource Description Returns a description of the specified resource in the specified stack. For deleted stacks, describe-stack-resource returns resource information for up to 90 days after the stack has been deleted.
Reference: http://docs.aws.amazon.com/cli/latest/reference/cloudformation/describe-stack-resource.html

QUESTION 537
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 two.)

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 538
A company has a three-tier application running on AWS with a web server, an application server, and an Amazon RDS MySQL DB instance. A solutions architect is designing a disaster recovery (DR) solution with an RPO of 5 minutes.
Which solution will meet the company's requirements?

A. Configure AWS Backup to perform cross-Region backups of all servers every 5 minutes. Reprovision the three tiers in the DR Region from the backups using AWS CloudFormation in the event of a disaster.
B. Maintain another running copy of the web and application server stack in the DR Region using AWS CloudFormation drift detection. Configure cross-Region snapshots of the DB instance to the DR Region every 5 minutes. In the event of a disaster, restore the DB instance using the snapshot in the DR Region.
C. Use Amazon EC2 Image Builder to create and copy AMIs of the web and application server to both the primary and DR Regions. Create a cross-Region read replica of the DB instance in the DR Region. In the event of a disaster, promote the read replica to become the master and reprovision the servers with AWS CloudFormation using the AMIs.
D. Create AMIs of the web and application servers in the DR Region. Use scheduled AWS Glue jobs to synchronize the DB instance with another DB instance in the DR Region. In the event of a disaster, switch to the DB instance in the DR Region and reprovision the servers with AWS CloudFormation using the AMIs.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 539
Which of the following statements is correct about AWS Direct Connect?

A. Connections to AWS Direct Connect require double clad fiber for 1 gigabit Ethernet with Auto Negotiation enabled for the port.

B. An AWS Direct Connect location provides access to Amazon Web Services in the region it is associated with.
C. AWS Direct Connect links your internal network to an AWS Direct Connect location over a standard 50 gigabit Ethernet cable.
D. To use AWS Direct Connect, your network must be collocated with a new AWS Direct Connect location.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
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. 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. To use AWS Direct Connect, your network is collocated with an existing AWS Direct Connect location. 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.
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.html

QUESTION 540
A large financial company is deploying applications that consist of Amazon EC2 and Amazon RDS instances to the AWS Cloud using AWS CloudFormation.
The CloudFormation stack has the following stack policy:


```
{
      "Statement" : [
         {
            "Effect" : "Allow",
            "Action" : ["Update:*"],
            "Principal": "*",
            "Resource" : "*"
         }
      ]
}
```

The company wants to ensure that developers do not lose data by accidentally removing or replacing RDS instances when updating the CloudFormation stack.
Developers also still need to be able to modify or remove EC2 instances as needed.
How should the company change the stack policy to meet these requirements?

A. Modify the statement to specify “Effect”: “Deny”, “Action”:[“Update:*”] for all logical RDS resources.
B. Modify the statement to specify “Effect”: “Deny”, “Action”:[“Update:Delete”] for all logical RDS resources.
C. Add a second statement that specifies “Effect”: “Deny”, “Action”:[“Update:Delete”, “Update:Replace”] for all logical RDS resources.
D. Add a second statement that specifies “Effect”: “Deny”, “Action”:[“Update:*”] for all logical RDS resources.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 541
A user has created a launch configuration for Auto Scaling where CloudWatch detailed monitoring is disabled. The user wants to now enable detailed monitoring. How can the user achieve this?

A. Update the Launch config with CLI to set InstanceMonitoringDisabled = false
B. The user should change the Auto Scaling group from the AWS console to enable detailed monitoring
C. Create a new Launch Config with detail monitoring enabled and update the Auto Scaling group
D. Update the Launch config with CLI to set InstanceMonitoring.Enabled = true

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
CloudWatch is used to monitor AWS as well as the custom services. To enable detailed instance monitoring for a new Auto Scaling group, the user does not need to take any extra steps. When the user creates the Auto Scaling launch config as the first step for creating an Auto Scaling group, each launch configuration contains a flag named InstanceMonitoring.Enabled.
The default value of this flag is true. When the user has created a launch configuration with InstanceMonitoring.Enabled = false it will involve multiple steps to enable detail monitoring. The steps are: Create a new Launch config with detailed monitoring enabled Update the Auto Scaling group with a new launch config Enable detail monitoring on each EC2 instance
Reference: http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/as-metricscollected.html

QUESTION 542
A company has a new application that needs to run on five Amazon EC2 instances in a single AWS Region. The application requires high-throughput, low-latency network connections between all of the EC2 instances where the application will run.
There is no requirement for the application to be fault tolerant.
Which solution will meet these requirements?

A. Launch five new EC2 instances into a cluster placement group. Ensure that the EC2 instance type supports enhanced networking.
B. Launch five new EC2 instances into an Auto Scaling group in the same Availability Zone. Attach an extra elastic network interface to each EC2 instance.
C. Launch five new EC2 instances into a partition placement group. Ensure that the EC2 instance type supports enhanced networking.

D. Launch five new EC2 instances into a spread placement group. Attach an extra elastic network interface to each EC2 instance.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 543
How much memory does the cr1.8xlarge instance type provide?

A. 224 GB
B. 124 GB
C. 184 GB
D. 244 GB

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The CR1 instances are part of the memory optimized instances. They offer lowest cost per GB RAM among all the AWS instance families. CR1 instances are part of the new generation of memory optimized instances, which can offer up to 244 GB RAM and run on faster CPUs (Intel Xeon E5-2670 with NUMA support) in comparison to the M2 instances of the same family. They support cluster networking for bandwidth intensive applications. cr1.8xlarge is one of the largest instance types of the CR1 family, which can offer 244 GB RAM.
Reference: http://aws.amazon.com/ec2/instance-types/

QUESTION 544
A company is finalizing the architecture for its backup solution for applications running on AWS. All of the applications run on AWS and use at least two Availability Zones in each tier.
Company policy requires IT to durably store nightly backups for all its data in at least two locations: production and disaster recovery. The locations must be in different geographic regions. The company also needs the backup to be available to restore immediately at the production data center, and within 24 hours at the disaster recovery location. All backup processes must be fully automated.
What is the MOST cost-effective backup solution that will meet all requirements?

A. Back up all the data to a large Amazon EBS volume attached to the backup media server in the production region. Run automated scripts to snapshot these volumes nightly, and copy these snapshots to the disaster recovery region.
B. Back up all the data to Amazon S3 in the disaster recovery region. Use a lifecycle policy to move this data to Amazon Glacier in the production region

immediately. Only the data is replicated; remove the data from the S3 bucket in the disaster recovery region.

C. Back up all the data to Amazon Glacier in the production region. Set up cross-region replication of this data to Amazon Glacier in the disaster recovery region. Set up a lifecycle policy to delete any data older than 60 days.

D. Back up all the data to Amazon S3 in the production region. Set up cross-region replication of this S3 bucket to another region and set up a lifecycle policy in the second region to immediately move this data to Amazon Glacier.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 545
A company wants to change its internal cloud billing strategy for each of its business units. Currently, the cloud governance team shares reports for overall cloud spending with the head of each business unit. The company uses AWS Organizations to manage the separate AWS accounts for each business unit. The existing tagging standard in Organizations includes the application, environment, and owner. The cloud governance team wants a centralized solution so each business unit receives monthly reports on its cloud spending. The solution should also send notifications for any cloud spending that exceeds a set threshold.
Which solution is the MOST cost-effective way to meet these requirements?

A. Configure AWS Budgets in each account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in each account to create monthly reports for each business unit.

B. Configure AWS Budgets in the organization's master account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in the organization's master account to create monthly reports for each business unit.

C. Configure AWS Budgets in each account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use the AWS Billing and Cost Management dashboard in each account to create monthly reports for each business unit.

D. Enable AWS Cost and Usage Reports in the organization's master account and configure reports grouped by application, environment, and owner. Create an AWS Lambda function that processes AWS Cost and Usage Reports, sends budget alerts, and sends monthly reports to each business unit's email list.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 546
A company has a web service deployed in the following two AWS Regions: us-west-2 and us-est-1. Each AWS region runs an identical version of the web service. Amazon Route 53 is used to route customers to the AWS Region that has the lowest latency.
The company wants to improve the availability of the web service in case an outage occurs in one of the two AWS Regions.

A Solutions Architect has recommended that a Route 53 health check be performed. The health check must detect a specific text on an endpoint. What combination of conditions should the endpoint meet to pass the Route 53 health check? (Choose two.)

A. The endpoint must establish a TCP connection within 10 seconds.
B. The endpoint must return an HTTP 200 status code.
C. The endpoint must return an HTTP 2xx or 3xx status code.
D. The specific text string must appear within the first 5,120 bytes of the response.
E. The endpoint must respond to the request within the number of seconds specified when creating the health check.

Correct Answer: CD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 547
An ecommerce website running on AWS uses an Amazon RDS for MySQL DB instance with General Purpose SSD storage.
The developers chose an appropriate instance type based on demand, and configured 100 GB of storage with a sufficient amount of free space.
The website was running smoothly for a few weeks until a marketing campaign launched. On the second day of the campaign, users reported long wait times and time outs. Amazon CloudWatch metrics indicated that both reads and writes to the DB instance were experiencing long response times. The CloudWatch metrics show 40% to 50% CPU and memory utilization, and sufficient free storage space is still available. The application server logs show no evidence of database connectivity issues.
What could be the root cause of the issue with the marketing campaign?

A. It exhausted the I/O credit balance due to provisioning low disk storage during the setup phase.
B. It caused the data in the tables to change frequently, requiring indexes to be rebuilt to optimize queries.
C. It exhausted the maximum number of allowed connections to the database instance.
D. It exhausted the network bandwidth available to the RDS for MySQL DB instance.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 548
AWS has launched T2 instances which come with CPU usage credit. An organization has a requirement which keeps an instance running for 24 hours. However,

the organization has high usage only during 11 AM to 12 PM. The organization is planning to use a T2 small instance for this purpose.
If the organization already has multiple instances running since Jan 2012, which of the below mentioned options should the organization implement while launching a T2 instance?

A. The organization must migrate to the EC2-VPC platform first before launching a T2 instance.
B. While launching a T2 instance the organization must create a new AWS account as this account does not have the EC2- VPC platform.
C. Create a VPC and launch a T2 instance as part of one of the subnets of that VPC.
D. While launching a T2 instance the organization must select EC2-VPC as the platform.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. The AWS account provides two platforms:
EC2-CLASSIC and EC2-VPC, depending on when the user has created his AWS account and which regions he is using. If the user has created the AWS account after 2013-12-04, it supports only EC2VPC. In this scenario, since the account is before the required date the supported platform will be EC2-CLASSIC. It is required that the organization creates a VPC as the T2 instances can be launched only as a part of VPC.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/vpc-migrate.html

QUESTION 549
A company has a website that enables users to upload videos. Company policy states the uploaded videos must be analyzed for restricted content. An uploaded video is placed in Amazon S3, and a message is pushed to an Amazon SQS queue with the video's location. A backend application pulls this location from Amazon SQS and analyzes the video.
The video analysis is compute-intensive and occurs sporadically during the day. The website scales with demand. The video analysis application runs on a fixed number of instances. Peak demand occurs during the holidays, so the company must add instances to the application during this time. All instances used are currently on-demand Amazon EC2 T2 instances.
The company wants to reduce the cost of the current solution.
Which of the following solutions is MOST cost-effective?

A. Keep the website on T2 instances. Determine the minimum number of website instances required during off-peak times and use Spot Instances to cover them while using Reserved Instances to cover peak demand. Use Amazon EC2 R4 and Amazon EC2 R5 Reserved Instances in an Auto Scaling group for the video analysis application.
B. Keep the website on T2 instances. Determine the minimum number of website instances required during off-peak times and use Reserved Instances to cover them while using On-Demand Instances to cover peak demand. Use Spot Fleet for the video analysis application comprised of Amazon EC2 C4 and Amazon EC2 C5 Spot Instances.
C. Migrate the website to AWS Elastic Beanstalk and Amazon EC2 C4 instances. Determine the minimum number of website instances required during off-peak times and use On-Demand Instances to cover them while using Spot capacity to cover peak demand. Use Spot Fleet for the video analysis application

comprised of C4 and Amazon EC2 C5 instances.

D. Migrate the website to AWS Elastic Beanstalk and Amazon EC2 R4 instances. Determine the minimum number of website instances required during off-peak times and use Reserved Instances to cover them while using On-Demand Instances to cover peak demand. Use Spot Fleet for the video analysis application comprised of R4 and Amazon EC2 R5 instances.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 550
As a part of building large applications in the AWS Cloud, the Solutions Architect is required to implement the perimeter security protection. Applications running on AWS have the following endpoints:
Application Load Balancer
Amazon API Gateway regional endpoint Elastic IP address-based EC2 instances.
Amazon S3 hosted websites. Classic Load Balancer
The Solutions Architect must design a solution to protect all of the listed web front ends and provide the following security capabilities:
DDoS protection
SQL injection protection
IP address whitelist/blacklist
HTTP flood protection
Bad bot scraper protection
How should the Solutions Architect design the solution?

A. Deploy AWS WAF and AWS Shield Advanced on all web endpoints. Add AWS WAF rules to enforce the company's requirements.

B. Deploy Amazon CloudFront in front of all the endpoints. The CloudFront distribution provides perimeter protection. Add AWS Lambda-based automation to provide additional security.

C. Deploy Amazon CloudFront in front of all the endpoints. Deploy AWS WAF and AWS Shield Advanced. Add AWS WAF rules to enforce the company's requirements. Use AWS Lambda to automate and enhance the security posture.

D. Secure the endpoints by using network ACLs and security groups and adding rules to enforce the company's requirements. Use AWS Lambda to automatically update the rules.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 551
A company has several teams, and each team has their own Amazon RDS database that totals 100 TB. The company is building a data query platform for Business Intelligence Analysts to generate a weekly business report. The new system must run ad-hoc SQL queries.
What is the MOST cost-effective solution?

A. Create a new Amazon Redshift cluster. Create an AWS Glue ETL job to copy data from the RDS databases to the Amazon Redshift cluster. Use Amazon Redshift to run the query.
B. Create an Amazon EMR cluster with enough core nodes. Run an Apache Spark job to copy data from the RDS databases to a Hadoop Distributed File System (HDFS). Use a local Apache Hive metastore to maintain the table definition. Use Spark SQL to run the query.
C. Use an AWS Glue ETL job to copy all the RDS databases to a single Amazon Aurora PostgreSQL database. Run SQL queries on the Aurora PostgreSQL database.
D. Use an AWS Glue crawler to crawl all the databases and create tables in the AWS Glue Data Catalog. Use an AWS Glue ETL job to load data from the RDS databases to Amazon S3, and use Amazon Athena to run the queries.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 552
A company needs to implement a patching process for its servers. The on-premises servers and Amazon EC2 instances use a variety of tools to perform patching. Management requires a single report showing the patch status of all the servers and instances.
Which set of actions should a solutions architect take to meet these requirements?

A. Use AWS Systems Manager to manage patches on the on-premises servers and EC2 instances. Use Systems Manager to generate patch compliance reports
B. Use AWS OpsWorks to manage patches on the on-premises servers and EC2 instances. Use Amazon QuickSight integration with OpsWorks to generate patch compliance reports.
C. Use an Amazon EventBridge (Amazon CloudWatch Events) rule to apply patches by scheduling an AWS Systems Manager patch remediation job. Use Amazon Inspector to generate patch compliance reports.
D. Use AWS OpsWorks to manage patches on the on-premises servers and EC2 instances. Use AWS X-Ray to post the patch status to AWS Systems Manager OpsCenter to generate patch compliance reports.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
You can use AWS Systems Manager Configuration Compliance to scan your fleet of managed instances for patch compliance.
Reference: https://aws.amazon.com/blogs/mt/how-moodys-uses-aws-systems-manager-to-patch-servers-across-multiplecloud-providers/

QUESTION 553
In Amazon VPC, what is the default maximum number of BGP advertised routes allowed per route table?

A. 15
B. 100
C. 5
D. 10

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The maximum number of BGP advertised routes allowed per route table is 100.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Appendix_Limits.html

QUESTION 554
Which statement is NOT true about a stack which has been created in a Virtual Private Cloud (VPC) in AWS OpsWorks?

A. Subnets whose instances cannot communicate with the Internet are referred to as public subnets.
B. Subnets whose instances can communicate only with other instances in the VPC and cannot communicate directly with the Internet are referred to as private subnets.
C. All instances in the stack should have access to any package repositories that your operating system depends on, such as the Amazon Linux or Ubuntu Linux repositories.
D. Your app and custom cookbook repositories should be accessible for all instances in the stack.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In AWS OpsWorks, you can control user access to a stack's instances by creating it in a virtual private cloud (VPC). For example, you might not want users to

have direct access to your stack's app servers or databases and instead require that all public traffic be channeled through an Elastic Load Balancer. A VPC consists of one or more subnets, each of which contains one or more instances. Each subnet has an associated routing table that directs outbound traffic based on its destination IP address. Instances within a VPC can generally communicate with each other, regardless of their subnet.
Subnets whose instances can communicate with the Internet are referred to as public subnets. Subnets whose instances can communicate only with other instances in the VPC and cannot communicate directly with the Internet are referred to as private subnets. AWS OpsWorks requires the VPC to be configured so that every instance in the stack, including instances in private subnets, has access to the following endpoints:
The AWS OpsWorks service, https://opsworks-instance-service.us-east-1.amazonaws.com.
Amazon S3
The package repositories for Amazon Linux or Ubuntu 12.04 LTS, depending on which operating system you specify. Your app and custom cookbook repositories.
Reference: http://docs.aws.amazon.com/opsworks/latest/userguide/workingstacks-vpc.html#workingstacks-vpc-basics

QUESTION 555
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 two.)

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 medisabled 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.

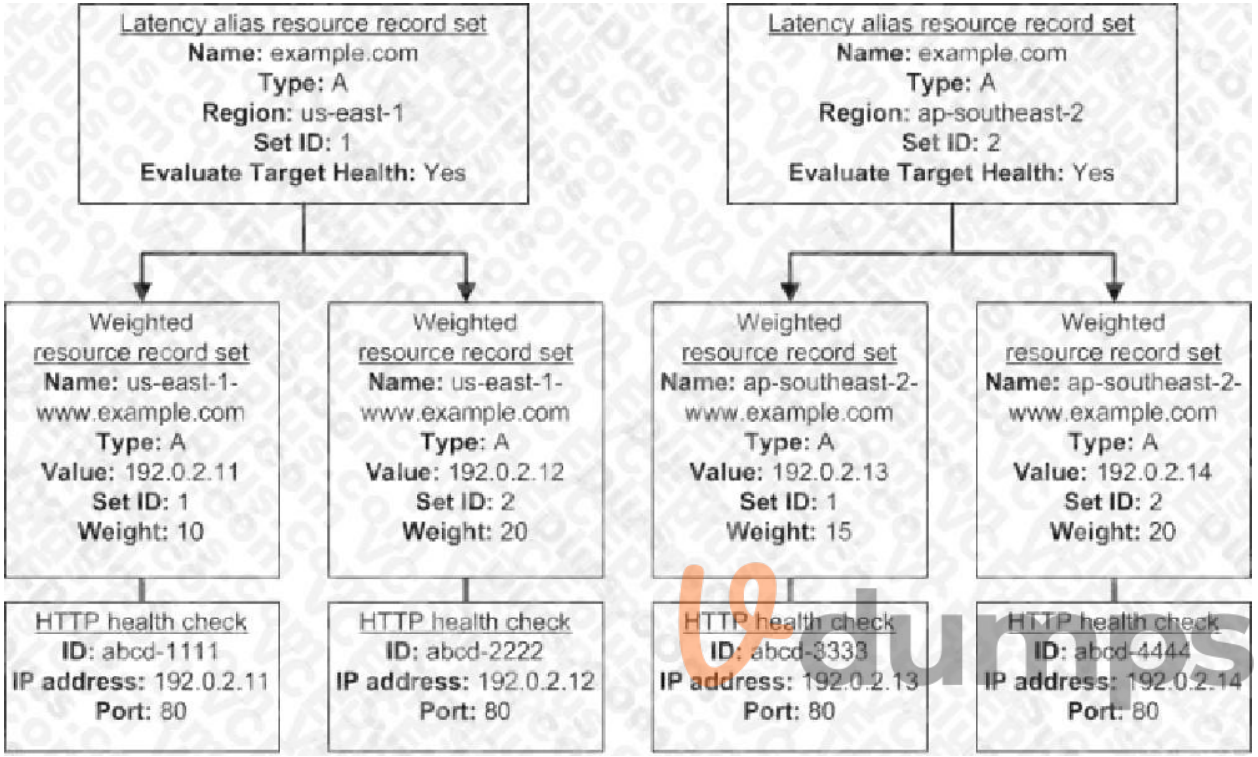

Here's how Amazon EC2 and Amazon Route 53 are configured:

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.

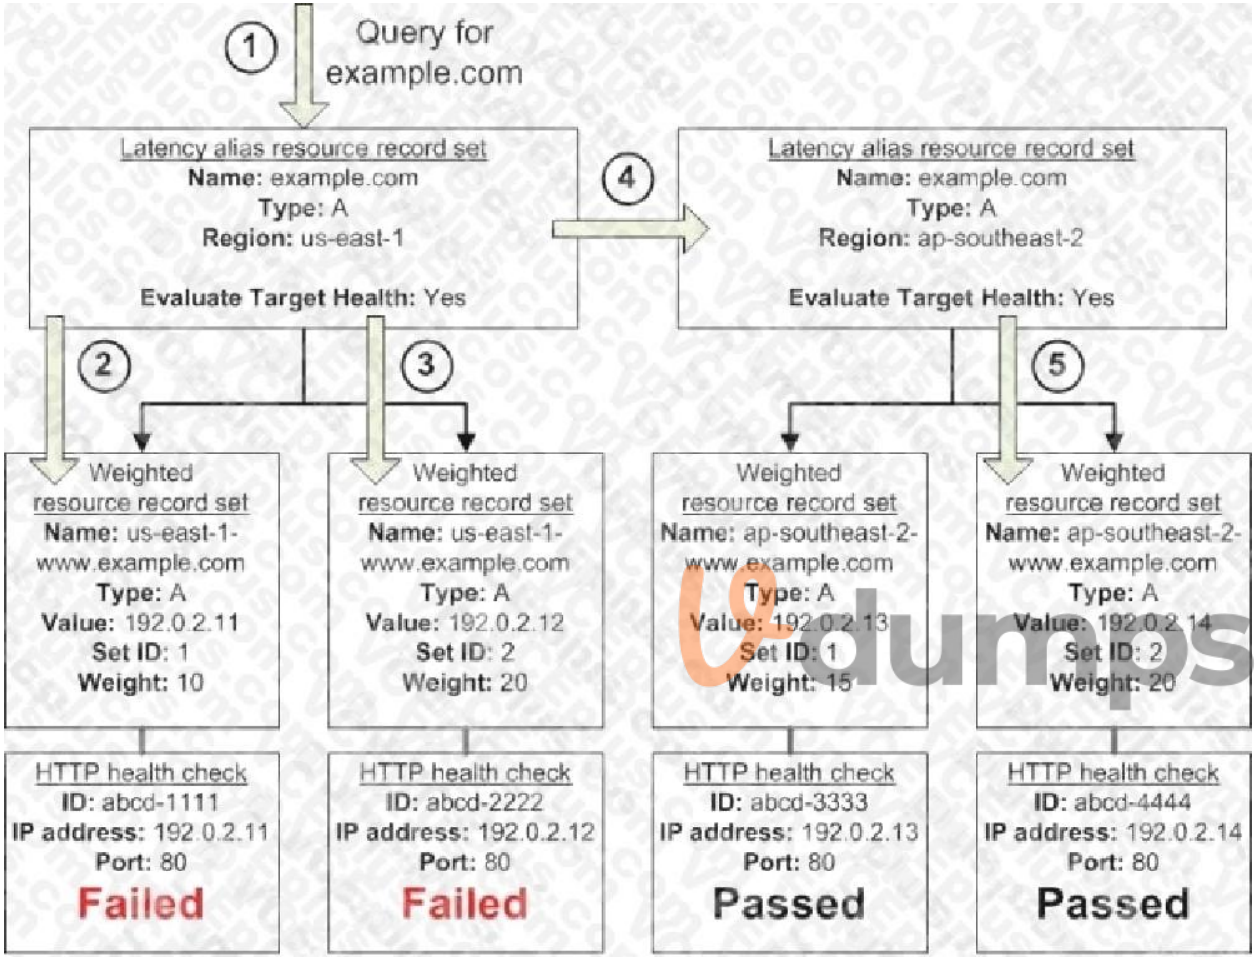

The preceding diagram illustrates the following sequence of events:

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.

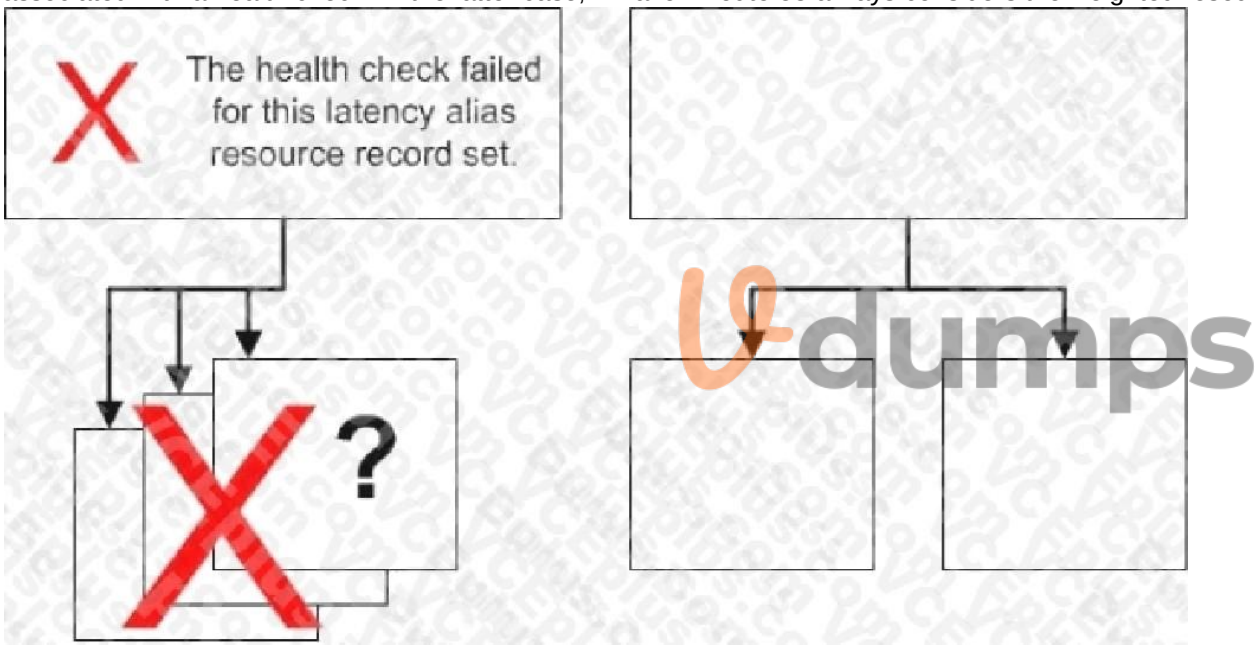

If the health check for the latency alias resource record set fails, Amazon Route 53 stops responding to queries using any of the weighted resource record sets in the alias target, even if they're all healthy. Amazon Route 53 doesn't know the status of the weighted resource record sets because it never 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:

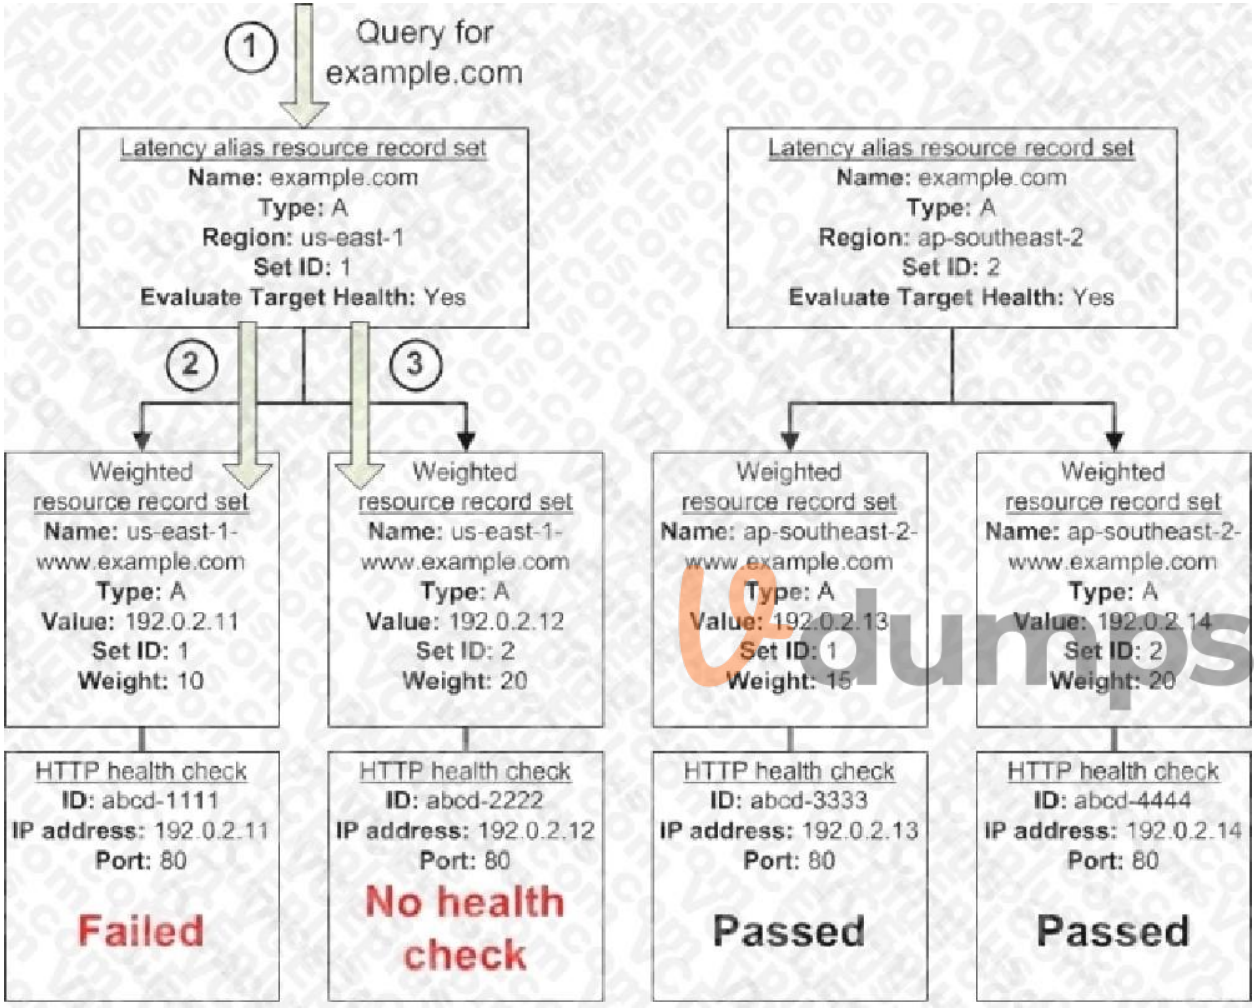

Here's what happens when you omit a health check on a non-alias resource record set in this configuration:

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:

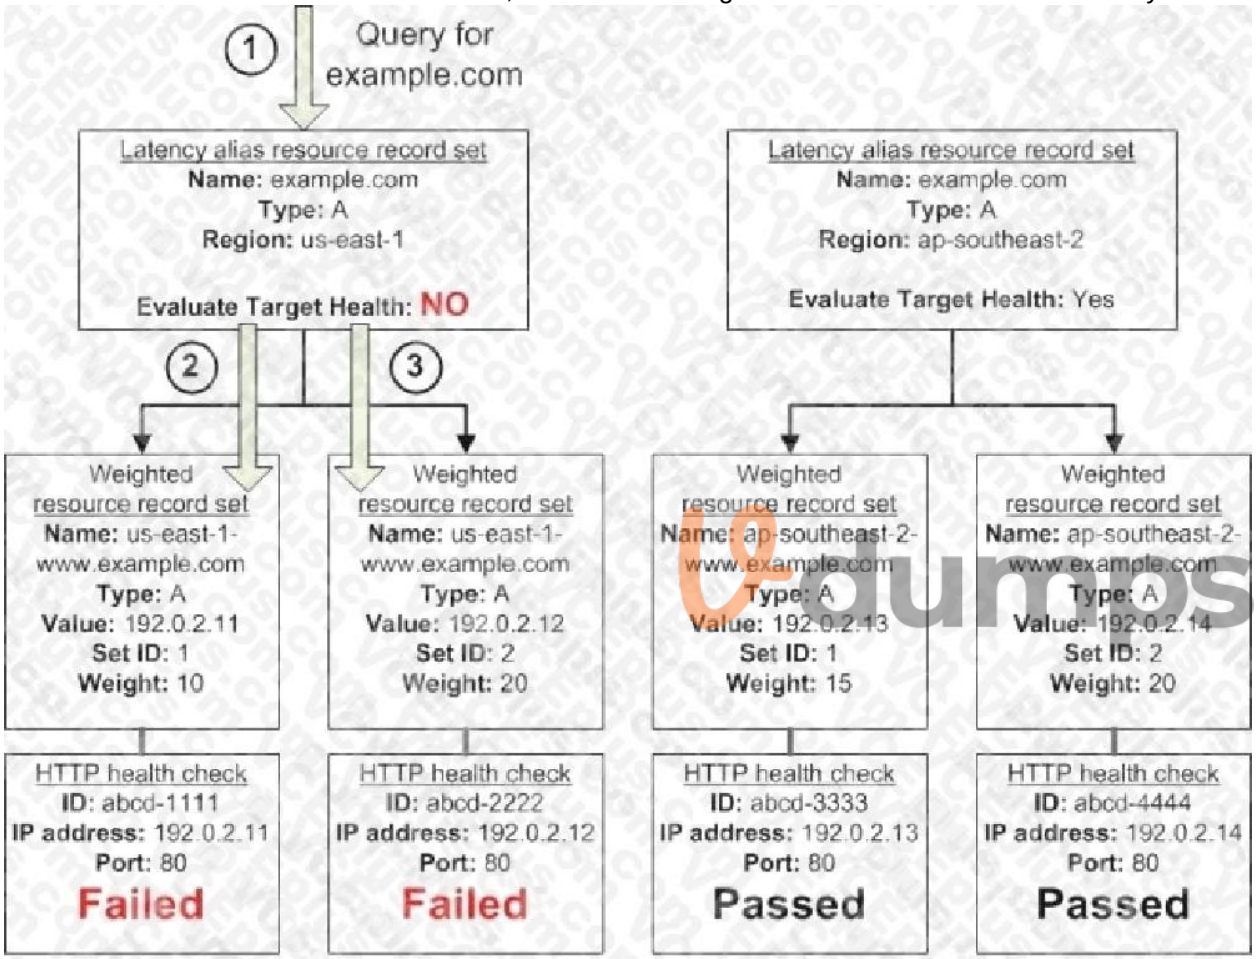

Here's what happens when you set Evaluate Target Health to No for an alias resource record set in this configuration:
Amazon Route 53 receives a query for example.com. Based on the latency for the user making the request, Amazon Route 53 selects the latency alias resource record set for the us-east-1 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 556

A company uses AWS Organizations to manage one parent account and nine member accounts. The number of member accounts is expected to grow as the business grows. A security engineer has requested consolidation of AWS CloudTrail logs into the parent account for compliance purposes. Existing logs currently stored in Amazon S3 buckets in each individual member account should not be lost.
Future member accounts should comply with the logging strategy.
Which operationally efficient solution meets these requirements?

A. Create an AWS Lambda function in each member account with a cross-account role. Trigger the Lambda functions when new CloudTrail logs are created and copy the CloudTrail logs to a centralized S3 bucket. Set up an Amazon CloudWatch alarm to alert if CloudTrail is not configured properly.
B. Configure CloudTrail in each member account to deliver log events to a central S3 bucket. Ensure the central S3 bucket policy allows PutObject access from the member accounts. Migrate existing logs to the central S3 bucket. Set up an Amazon CloudWatch alarm to alert if CloudTrail is not configured properly.
C. Configure an organization-level CloudTrail in the parent account to deliver log events to a central S3 bucket. Migrate the existing CloudTrail logs from each member account to the central S3 bucket. Delete the existing CloudTrail and logs in the member accounts.
D. Configure an organization-level CloudTrail in the parent account to deliver log events to a central S3 bucket. Configure CloudTrail in each member account to deliver log events to the central S3 bucket.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/blogs/architecture/stream-amazon-cloudwatch-logs-to-a-centralized-account-for-auditand-analysis/

QUESTION 557
A company is running a distributed application on a set of Amazon EC2 instances in an Auto Scaling group. The application stores large amounts of data on an Amazon Elastic File System (Amazon EFS) file system, and new data is generated monthly. The company needs to back up the data in a secondary AWS Region to restore from in case of a performance problem in its primary Region. The company has an RTO of 1 hour. A solutions architect needs to create a backup strategy while minimizing the extra cost.
Which backup strategy should the solutions architect recommend to meet these requirements?

A. Create a pipeline in AWS Data Pipeline. Copy the data to an EFS file system in the secondary Region. Create a lifecycle policy to move files to the EFS One Zone-Infrequent Access storage class.
B. Set up automatic backups by using AWS Backup. Create a copy rule to copy backups to an Amazon S3 bucket in the secondary Region. Create a lifecycle policy to move backups to the S3 Glacier storage class.
C. Set up AWS DataSync and continuously copy the files to an Amazon S3 bucket in the secondary Region. Create a lifecycle policy to move files to the S3 Glacier Deep Archive storage class.
D. Turn on EFS Cross-Region Replication and set the secondary Region as the target. Create a lifecycle policy to move files to the EFS Infrequent Access storage class in the secondary Region.

Correct Answer: A

Section: (none)
Explanation

Explanation/Reference:

QUESTION 558
A company plans to migrate to AWS. A solutions architect uses AWS Application Discovery Service over the fleet and discovers that there is an Oracle data warehouse and several PostgreSQL databases. Which combination of migration patterns will reduce licensing costs and operational overhead? (Choose two.)

A. Lift and shift the Oracle data warehouse to Amazon EC2 using AWS DMS.
B. Migrate the Oracle data warehouse to Amazon Redshift using AWS SCT and AWS DMS
C. Lift and shift the PostgreSQL databases to Amazon EC2 using AWS DMS.
D. Migrate the PostgreSQL databases to Amazon RDS for PostgreSQL using AWS DMS.
E. Migrate the Oracle data warehouse to an Amazon EMR managed cluster using AWS DMS.

Correct Answer: DE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 559
In CloudFormation, if you want to map an Amazon Elastic Block Store to an Amazon EC2 instance, _____.

A. you reference the logical IDs to associate the block stores with the instance
B. you reference the physical IDs of the instance along with the resource type
C. you reference the instance IDs of the block store along with the resource properties
D. you reference the physical IDs of both the block stores and the instance

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In AWS CloudFormation, if you want to map an Amazon Elastic Block Store to an Amazon EC2 instance, you reference the logical IDs to associate the block stores with the instance.

Reference: http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/concept-resources.html

QUESTION 560
A company is developing a gene reporting device that will collect genomic information to assist researchers will collecting large samples of data from a diverse population. The device will push 8 KB of genomic data every second to a data platform that will need to process and analyze the data and provide information back to researchers. The data platform must meet the following requirements:
Provide near-real-time analytics of the inbound genomic data
Ensure the data is flexible, parallel, and durable
Deliver results of processing to a data warehouse
Which strategy should a solutions architect use to meet these requirements?

A. Use Amazon Kinesis Data Firehouse to collect the inbound sensor data, analyze the data with Kinesis clients, and save the results to an Amazon RDS instance.
B. Use Amazon Kinesis Data Streams to collect the inbound sensor data, analyze the data with Kinesis clients, and save the results to an Amazon Redshift cluster using Amazon EMR.
C. Use Amazon S3 to collect the inbound device data, analyze the data from Amazon SQS with Kinesis, and save the results to an Amazon Redshift cluster.
D. Use an Amazon API Gateway to put requests into an Amazon SQS queue, analyze the data with an AWS Lambda function, and save the results to an Amazon Redshift cluster using Amazon EMR.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 561
A company runs an application on AWS. An AWS Lambda function uses credentials to authenticate to an Amazon RDS for MySQL DB instance. A security risk assessment identified that these credentials are not frequently rotated. Also, encryption at rest is not enabled for the DB instance. The security team requires that both of these issues be resolved.
Which strategy should a solutions architect recommend to remediate these security risks?

A. Configure the Lambda function to store and retrieve the database credentials in AWS Secrets Manager and enable rotation of the credentials. Take a snapshot of the DB instance and encrypt a copy of that snapshot. Replace the DB instance with a new DB instance that is based on the encrypted snapshot.
B. Enable IAM DB authentication on the DB instance. Grant the Lambda execution role access to the DB instance. Modify the DB instance and enable encryption.
C. Enable IAM DB authentication on the DB instance. Grant the Lambda execution role access to the DB instance. Create an encrypted read replica of the DB instance. Promote the encrypted read replica to be the new primary node.
D. Configure the Lambda function to store and retrieve the database credentials as encrypted AWS Systems Manager Parameter Store parameters. Create another Lambda function to automatically rotate the credentials. Create an encrypted read replica of the DB instance. Promote the encrypted read replica to

be the new primary node.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/secretsmanager/latest/userguide/enable-rotation-rds.html

QUESTION 562
A company has implemented AWS Organizations. It has recently set up a number of new accounts and wants to deny access to a specific set of AWS services in these new accounts.
How can this be controlled MOST efficiently?

A. Create an IAM policy in each account that denies access to the services. Associate the policy with an IAM group, and add all IAM users to the group.
B. Create a service control policy that denies access to the services. Add all of the new accounts to a single organizational unit (OU), and apply the policy to that OU.
C. Create an IAM policy in each account that denies access to the services. Associate the policy with an IAM role, and instruct users to log in using their corporate credentials and assume the IAM role.
D. Create a service control policy that denies access to the services, and apply the policy to the root of the organization.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp.html

QUESTION 563
In Amazon Redshift, how many slices does a dw2.8xlarge node have?

A. 16
B. 8
C. 32
D. 2

Correct Answer: C

Section: (none)
Explanation

Explanation/Reference:
Explanation:
The disk storage for a compute node in Amazon Redshift is divided into a number of slices, equal to the number of processor cores on the node. For example, each DW1.XL compute node has two slices, and each DW2.8XL compute node has 32 slices.
Reference: http://docs.aws.amazon.com/redshift/latest/dg/t_Distributing_data.html

QUESTION 564
When configuring your customer gateway to connect to your VPC, the_____Association is established first between the virtual private gateway and customer gateway using the Pre-Shared Key as the authenticator.

A. IPsec
B. BGP
C. IKE Security
D. Tunnel

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
When configuring your customer gateway to connect to your VPC, several steps need to be completed. The IKE Security Association is established first between the virtual private gateway and customer gateway using the Pre-Shared Key as the authenticator.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/NetworkAdminGuide/Introduction.html

QUESTION 565
A company is using multiple AWS accounts. The DNS records are stored in a private hosted zone for Amazon Route 53 in Account A. The company's applications and databases are running in Account B.
A solutions architect will deploy a two-tier application in a new VPC. To simplify the configuration, the db.example.com CNAME record set for the Amazon RDS endpoint was created in a private hosted zone for Amazon Route 53.
During deployment the application failed to start. Troubleshooting revealed that db.example.com is not resolvable on the Amazon EC2 instance. The solutions architect confirmed that the record set was created correctly in Route 53.
Which combination of steps should the solutions architect take to resolve this issue? (Choose two.)

A. Deploy the database on a separate EC2 instance in the new VPC. Create a record set for the instance's private IP in the private hosted zone. Delete the association authorization in Account A.
B. Use SSH to connect to the application tier EC2 instance. Add an RDS endpoint IP address to the /etc/resolv conf file.

Configure Route 53 replication between AWS accounts.

C. Create an authorization to associate the private hosted zone in Account A with the new VPC in Account B.

D. Create a private hosted zone for the example com domain in Account

E. Associate a new VPC in Account B with a hosted zone in Account

Correct Answer: BE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 566
A user is trying to create a vault in AWS Glacier. The user wants to enable notifications.
In which of the below mentioned options can the user enable the notifications from the AWS console?

A. Glacier does not support the AWS console

B. Archival Upload Complete

C. Vault Upload Job Complete

D. Vault Inventory Retrieval Job Complete

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
From AWS console the user can configure to have notifications sent to Amazon Simple Notifications Service (SNS). The user can select specific jobs that, on completion, will trigger the notifications such as Vault Inventory Retrieval Job Complete and Archive Retrieval Job Complete.
Reference: http://docs.aws.amazon.com/amazonglacier/latest/dev/configuring-notifications-console.html

QUESTION 567
AWS Direct Connect itself has NO specific resources for you to control access to. Therefore, there are no AWS Direct Connect Amazon Resource Names (ARNs) for you to use in an Identity and Access Management (IAM) policy.
With that in mind, how is it possible to write a policy to control access to AWS Direct Connect actions?

A. You can leave the resource name field blank.

B. You can choose the name of the AWS Direct Connection as the resource.

C. You can use an asterisk (*) as the resource.

D. You can create a name for the resource.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
AWS Direct Connect itself has no specific resources for you to control access to. Therefore, there are no AWS Direct Connect ARNs for you to use in an IAM policy. You use an asterisk (*) as the resource when writing a policy to control access to AWS Direct Connect actions.
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/using_iam.html

QUESTION 568
A user has configured two security groups which allow traffic as given below: 1: SecGrp1: Inbound on port 80 for 0.0.0.0/0 Inbound on port 22 for 0.0.0.0/0 2: SecGrp2:
Inbound on port 22 for 10.10.10.1/32
If both the security groups are associated with the same instance, which of the below mentioned statements is true?

A. It is not possible to have more than one security group assigned to a single instance
B. It is not possible to create the security group with conflicting rules. AWS will reject the request
C. It allows inbound traffic for everyone on both ports 22 and 80
D. It allows inbound traffic on port 22 for IP 10.10.10.1 and for everyone else on port 80

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A user can attach more than one security group to a single EC2 instance. In this case, the rules from each security group are effectively aggregated to create one set of rules. AWS uses this set of rules to determine whether to allow access or not.
Thus, here the rule for port 22 with IP 10.10.10.1/32 will merge with IP 0.0.0.0/0 and open ports 22 and 80 for all.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html

QUESTION 569
A company wants to allow its Marketing team to perform SQL queries on customer records to identify market segments. The data is spread across hundreds of files. The records must be encrypted in transit and at rest. The Team Manager must have the ability to manage users and groups, but no team members should have access to services or resources not required for the SQL queries. Additionally, Administrators need to audit the queries made and receive notifications when a query violates rules defined by the Security team.
AWS Organizations has been used to create a new account and an AWS IAM user with administrator permissions for the Team Manager.

Which design meets these requirements?

A. Apply a service control policy (SCP) that allows access to IAM, Amazon RDS, and AWS CloudTrail. Load customer records in Amazon RDS MySQL and train users to execute queries using the AWS CLI. Stream the query logs to Amazon CloudWatch Logs from the RDS database instance. Use a subscription filter with AWS Lambda functions to audit and alarm on queries against personal data.
B. Apply a service control policy (SCP) that denies access to all services except IAM, Amazon Athena, Amazon S3, and AWS CloudTrail. Store customer record files in Amazon S3 and train users to execute queries using the CLI via Athena.
Analyze CloudTrail events to audit and alarm on queries against personal data.
C. Apply a service control policy (SCP) that denies access to all services except IAM, Amazon DynamoDB, and AWS CloudTrail. Store customer records in DynamoDB and train users to execute queries using the AWS CLI. Enable DynamoDB streams to track the queries that are issued and use an AWS Lambda function for real-time monitoring and alerting.
D. Apply a service control policy (SCP) that allows access to IAM, Amazon Athena, Amazon S3, and AWS CloudTrail. Store customer records as files in Amazon S3 and train users to leverage the Amazon S3 Select feature and execute queries using the AWS CLI. Enable S3 object-level logging and analyze CloudTrail events to audit and alarm on queries against personal data.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 570
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 571
What feature of the load balancing service attempts to force subsequent connections to a service to be redirected to the same node as long as it is online?

A. Node balance
B. Session retention
C. Session multiplexing
D. Session persistence

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Session persistence is a feature of the load balancing service. It attempts to force subsequent connections to a service to be redirected to the same node as long as it is online.
Reference: http://docs.rackspace.com/loadbalancers/api/v1.0/clb-devguide/content/Concepts-d1e233.html

QUESTION 572
An internal security audit of AWS resources within a company found that a number of Amazon EC2 instances running Microsoft Windows workloads were missing several important operating system-level patches. A Solutions Architect has been asked to fix existing patch deficiencies, and to develop a workflow to ensure that future patching requirements are identified and taken care of quickly. The Solutions Architect has decided to use AWS Systems Manager. It is important that EC2 instance reboots do not occur at the same time on all Windows workloads to meet organizational uptime requirements.
Which workflow will meet these requirements in an automated manner?

A. Add a Patch Group tag with a value of Windows Servers to all existing EC2 instances. Ensure that all Windows EC2 instances are assigned this tag. Associate the AWS-DefaultPatchBaseline to the Windows Servers patch group. Define an AWS Systems Manager maintenance window, conduct patching within it, and associate it with the Windows Servers patch group. Register instances with the maintenance window using associated subnet IDs. Assign the AWS-RunPatchBaseline document as a task within each maintenance window.

B. Add a Patch Group tag with a value of Windows Servers to all existing EC2 instances. Ensure that all Windows EC2 instances are assigned this tag. Associate the AWS-WindowsPatchBaseline to the Windows Servers patch group. Create an Amazon CloudWatch Events rule configured to use a cron expression to schedule the execution of patching using the AWS Systems Manager run command. Assign the AWS-RunWindowsPatchBaseline document as a task associated with the Windows Servers patch group. Create an AWS Systems Manager State Manager document to define commands to be executed during patch execution.

C. Add a Patch Group tag with a value of either Windows Servers1 or Windows Servers2 to all existing EC2 instances. Ensure that all Windows EC2 instances are assigned this tag. Associate the AWSDefaultPatchBaseline with both Windows Servers patch groups. Define two non-overlapping AWS Systems Manager maintenance windows, conduct patching within them, and associate each with a different patch group. Register targets with specific maintenance windows using the Patch Group tags. Assign the AWS-RunPatchBaseline document as a task within each maintenance window.

D. Add a Patch Group tag with a value of either Windows Servers1 or Windows Server2 to all existing EC2 instances. Ensure that all Windows EC2 instances are assigned this tag. Associate the AWSWindowsPatchBaseline with both Windows Servers patch groups. Define two non-overlapping AWS Systems Manager maintenance windows, conduct patching within them, and associate each with a different patch group. Assign the AWS-RunWindowsPatchBaseline document as a task within each maintenance window. Create an AWS Systems Manager State Manager document to define commands to be executed during patch execution.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 573
A user is creating a PIOPS volume. What is the maximum ratio the user should configure between PIOPS and the volume size?

A. 5
B. 10
C. 20
D. 30

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Provisioned IOPS volumes are designed to meet the needs of I/O-intensive workloads, particularly database workloads that are sensitive to storage performance and consistency in random access I/O throughput. A provisioned IOPS volume can range in size from 10 GB to 1 TB and the user can provision up to 4000 IOPS per volume.

The ratio of IOPS provisioned to the volume size requested can be a maximum of 30; for example, a volume with 3000 IOPS must be at least 100 GB.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html

QUESTION 574
A company's AWS architecture currently uses access keys and secret access keys stored on each instance to access AWS services. Database credentials are hard-coded on each instance. SSH keys for command-line remote access are stored in a secured Amazon S3 bucket. The company has asked its solutions architect to improve the security posture of the architecture without adding operational complexity.
Which combination of steps should the solutions architect take to accomplish this? (Choose three.)

A. Use Amazon EC2 instance profiles with an IAM role
B. Use AWS Secrets Manager to store access keys and secret access keys
C. Use AWS Systems Manager Parameter Store to store database credentials
D. Use a secure fleet of Amazon EC2 bastion hosts for remote access
E. Use AWS KMS to store database credentials
F. Use AWS Systems Manager Session Manager for remote access

Correct Answer: ABD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 575
An organization is planning to extend their data center by connecting their DC with the AWS VPC using the VPN gateway.
The organization is setting up a dynamically routed VPN connection.
Which of the below mentioned answers is not required to setup this configuration?

A. The type of customer gateway, such as Cisco ASA, Juniper J-Series, Juniper SSG, Yamaha.
B. Elastic IP ranges that the organization wants to advertise over the VPN connection to the VPC.
C. Internet-routable IP address (static) of the customer gateway's external interface.
D. Border Gateway Protocol (BGP) Autonomous System Number (ASN) of the customer gateway.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:

The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. The organization wants to extend their network into the cloud and also directly access the internet from their AWS VPC. Thus, the organization should setup a Virtual Private Cloud (VPC) with a public subnet and a private subnet, and a virtual private gateway to enable communication with their data center network over an IPsec VPN tunnel. To setup this configuration the organization needs to use the Amazon VPC with a VPN connection. The organization network administrator must designate a physical appliance as a customer gateway and configure it.
The organization would need the below mentioned information to setup this configuration:
The type of customer gateway, such as Cisco ASA, Juniper J-Series, Juniper SSG, Yamaha Internet-routable IP address (static) of the customer gateway's external interface Border Gateway Protocol (BGP) Autonomous System Number (ASN) of the customer gateway, if the organization is creating a dynamically routed VPN connection. Internal network IP ranges that the user wants to advertise over the VPN connection to the VPC.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_VPN.html

QUESTION 576
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 577
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 two.)

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: AD
Section: (none)
Explanation

Explanation/Reference:
Explanation:
EC2 does not allow promiscuous mode, and you cannot put something in between the ELB and the web server (like a listener or IDP)

QUESTION 578
In which step of using AWS Direct Connect should the user determine the required port speed?

A. Complete the Cross Connect
B. Verify Your Virtual Interface
C. Download Router Configuration
D. Submit AWS Direct Connect Connection Request

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
To submit an AWS Direct Connect connection request, you need to provide the following information: Your contact information.
The AWS Direct Connect Location to connect to.
Details of AWS Direct Connect partner if you use the AWS Partner Network (APN) service. The port speed you require, either 1 Gbps or 10 Gbps.
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/getstarted.html#ConnectionRequest

QUESTION 579
A company operating a website on AWS requires high levels of scalability, availability, and performance. The company is running a Ruby on Rails application on Amazon EC2. It has a data tier on MySQL 5.6 on Amazon EC2 using 16 TB of Amazon EBS storage Amazon CloudFront is used to cache application content. The Operations team is reporting continuous and unexpected growth of EBS volumes assigned to the MySQL database. The Solutions Architect has been asked to design a highly scalable, highly available, and high-performing solution.
Which solution is the MOST cost-effective at scale?

A. Implement Multi-AZ and Auto Scaling for all EC2 instances in the current configuration. Ensure that all EC2 instances are purchased as reserved instances. Implement new elastic Amazon EBS volumes for the data tier.

B. Design and implement the Docker-based containerized solution for the application using Amazon ECS. Migrate to an Amazon Aurora MySQL Multi-AZ cluster. Implement storage checks for Aurora MySQL storage utilization and an AWS Lambda function to grow the Aurora MySQL storage, as necessary. Ensure that Multi-AZ architectures are implemented.

C. Ensure that EC2 instances are right-sized and behind an Elastic Load Balancing load balancer. Implement Auto Scaling with EC2 instances. Ensure that the reserved instances are purchased for fixed capacity and that Auto Scaling instances run on demand. Migrate to an Amazon Aurora MySQL Multi-AZ cluster. Ensure that Multi-AZ architectures are implemented.

D. Ensure that EC2 instances are right-sized and behind an Elastic Load Balancer. Implement Auto Scaling with EC2 instances. Ensure that Reserved instances are purchased for fixed capacity and that Auto Scaling instances run on demand.
Migrate to an Amazon Aurora MySQL Multi-AZ cluster. Implement storage checks for Aurora MySQL storage utilization and an AWS Lambda function to grow Aurora MySQL storage, as necessary. Ensure Multi-AZ architectures are implemented.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 580
The Principal element of an IAM policy refers to the specific entity that should be allowed or denied permission, whereas the translates to everyone except the specified entity.

A. NotPrincipal

B. Vendor

C. Principal

D. Action

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The element NotPrincipal that is included within your IAM policy statements allows you to specify an exception to a list of principals to whom the access to a specific resource is either allowed or denied. Use the NotPrincipal element to specify an exception to a list of principals. For example, you can deny access to all principals except the one named in the NotPrincipal element.

Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements.html#Principal

QUESTION 581
A company has multiple AWS accounts. The company recently had a security audit that revealed many unencrypted Amazon Elastic Block Store (Amazon EBS) volumes attached to Amazon EC2 instances.
A solutions architect must encrypt the unencrypted volumes and ensure that unencrypted volumes will be detected automatically in the future. Additionally, the company wants a solution that can centrally manage multiple AWS accounts with a focus on compliance and security.
Which combination of steps should the solutions architect take to meet these requirements? (Choose two.)

A. Create an organization in AWS Organizations. Set up AWS Control Tower, and turn on the strongly recommended guardrails. Join all accounts to the organization. Categorize the AWS accounts into OUs.
B. Use the AWS CLI to list all the unencrypted volumes in all the AWS accounts. Run a script to encrypt all the unencrypted volumes in place.
C. Create a snapshot of each unencrypted volume. Create a new encrypted volume from the unencrypted snapshot. Detach the existing volume, and replace it with the encrypted volume.
D. Create an organization in AWS Organizations. Set up AWS Control Tower, and turn on the mandatory guardrails. Join all accounts to the organization. Categorize the AWS accounts into OUs.
E. Turn on AWS CloudTrail. Configure an Amazon EventBridge (Amazon CloudWatch Events) rule to detect and automatically encrypt unencrypted volumes.

Correct Answer: AC
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/controltower/latest/userguide/guardrails.html
https://aws.amazon.com/premiumsupport/knowledge-center/create-unencrypted-volume-kms-key/

QUESTION 582
An organization has hosted an application on the EC2 instances. There will be multiple users connecting to the instance for setup and configuration of application. The organization is planning to implement certain security best practices.
Which of the below mentioned pointers will not help the organization achieve better security arrangement?

A. Allow only IAM users to connect with the EC2 instances with their own secret access key.
B. Create a procedure to revoke the access rights of the individual user when they are not required to connect to EC2 instance anymore for the purpose of application configuration.
C. Apply the latest patch of OS and always keep it updated.
D. Disable the password based login for all the users. All the users should use their own keys to connect with the instance securely.

Correct Answer: A

Section: (none)
Explanation

Explanation/Reference:
Explanation:
Since AWS is a public cloud any application hosted on EC2 is prone to hacker attacks. It becomes extremely important for a user to setup a proper security mechanism on the EC2 instances. A few of the security measures are listed below:
Always keep the OS updated with the latest patch
Always create separate users with in OS if they need to connect with the EC2 instances, create their keys and disable their password Create a procedure using which the admin can revoke the access of the user when the business work on the EC2 instance is completed. Lock down unnecessary ports.
Audit any proprietary applications that the user may be running on the EC2 instance Provide temporary escalated privileges, such as sudo for users who need to perform occasional privileged tasks The IAM is useful when users are required to work with AWS resources and actions, such as launching an instance. It is not useful to connect (RDP / SSH) with an instance.
Reference: http://aws.amazon.com/articles/1233/

QUESTION 583
You have custom Network File System (NFS) client settings for your Amazon Elastic File System (EFS). It takes up to three seconds for an Amazon Elastic Compute Cloud (EC2) instance to see a write operation performed on a file system from another Amazon EC2 instance.
Which of the following actions should you take to solve the custom NFS settings from causing delays in the write operation?

A. Unmount and remount the file system with the noac option to disable attribute caching.
B. Reduce the number of active users that have files open simultaneously on the instances.
C. Verify that the IP address of the specified mount target is valid.
D. Run the write operation from a different user ID on the same Amazon EC2 instance.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
When you set up custom NFS client settings, it takes up to three seconds for an Amazon EC2 instance to see a write operation being performed on a file system from another Amazon EC2 instance. To solve this issue, you must unmount and remount your file system with the noac option to disable attribute caching if the NFS client on the Amazon EC2 instance that is reading the data has attribute caching activated. Attribute cache can also be cleared on demand by using a programming language that is compatible with the NFS procedures. To do this, you must send an ACCESS procedure request immediately before a read request.
Reference: http://docs.aws.amazon.com/efs/latest/ug/troubleshooting.html#custom-nfs-settings-write-delays

QUESTION 584
A company has an environment that has a single AWS account. A solutions architect is reviewing the environment to recommend what the company could

improve specifically in terms of access to the AWS Management Console. The company's IT support workers currently access the console for administrative tasks, authenticating with named IAM users that have been mapped to their job role.
The IT support workers no longer want to maintain both their Active Directory and IAM user accounts. They want to be able to access the console by using their existing Active Directory credentials. The solutions architect is using AWS Single Sign- On (AWS SSO) to implement this functionality.
Which solution will meet these requirements MOST cost-effectively?

A. Create an organization in AWS Organizations. Turn on the AWS SSO feature in Organizations Create and configure a directory in AWS Directory Service for Microsoft Active Directory (AWS Managed Microsoft AD) with a two-way trust to the company's on-premises Active Directory. Configure AWS SSO and set the AWS Managed Microsoft AD directory as the identity source. Create permission sets and map them to the existing groups within the AWS Managed Microsoft AD directory.
B. Create an organization in AWS Organizations. Turn on the AWS SSO feature in Organizations Create and configure an AD Connector to connect to the company's on-premises Active Directory. Configure AWS SSO and select the AD Connector as the identity source. Create permission sets and map them to the existing groups within the company's Active Directory.
C. Create an organization in AWS Organizations. Turn on all features for the organization. Create and configure a directory in AWS Directory Service for Microsoft Active Directory (AWS Managed Microsoft AD) with a two-way trust to the company's on-premises Active Directory. Configure AWS SSO and select the AWS Managed Microsoft AD directory as the identity source. Create permission sets and map them to the existing groups within the AWS Managed Microsoft AD directory.
D. Create an organization in AWS Organizations. Turn on all features for the organization. Create and configure an AD Connector to connect to the company's on-premises Active Directory. Configure AWS SSO and select the AD Connector as the identity source. Create permission sets and map them to the existing groups within the company's Active Directory.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/single-sign-on/faqs/

QUESTION 585
A user authenticating with Amazon Cognito will go through a multi-step process to bootstrap their credentials. Amazon Cognito has two different flows for authentication with public providers.
Which of the following are the two flows?

A. Authenticated and non-authenticated
B. Public and private
C. Enhanced and basic
D. Single step and multistep

Correct Answer: C

Section: (none)
Explanation

Explanation/Reference:
Explanation:
A user authenticating with Amazon Cognito will go through a multi-step process to bootstrap their credentials. Amazon Cognito has two different flows for authentication with public providers: enhanced and basic.
Reference: http://docs.aws.amazon.com/cognito/devguide/identity/concepts/authentication-flow/

QUESTION 586
A public retail web application uses an Application Load Balancer (ALB) in front of Amazon EC2 instances running across multiple Availability Zones (AZs) in a Region backed by an Amazon RDS MySQL Multi-AZ deployment. Target group health checks are configured to use HTTP and pointed at the product catalog page. Auto Scaling is configured to maintain the web fleet size based on the ALB health check.
Recently, the application experienced an outage. Auto Scaling continuously replaced the instances during the outage. A subsequent investigation determined that the web server metrics were within the normal range, but the database tier was experiencing high load, resulting in severely elevated query response times.
Which of the following changes together would remediate these issues while improving monitoring capabilities for the availability and functionality of the entire application stack for future growth? (Choose two.)

A. Configure read replicas for Amazon RDS MySQL and use the single reader endpoint in the web application to reduce the load on the backend database tier.
B. Configure the target group health check to point at a simple HTML page instead of a product catalog page and the Amazon Route 53 health check against the product page to evaluate full application functionality. Configure Amazon CloudWatch alarms to notify administrators when the site fails.
C. Configure the target group health check to use a TCP check of the Amazon EC2 web server and the Amazon Route 53 health check against the product page to evaluate full application functionality. Configure Amazon CloudWatch alarms to notify administrators when the site fails.
D. Configure an Amazon CloudWatch alarm for Amazon RDS with an action to recover a high-load, impaired RDS instance in the database tier.
E. Configure an Amazon ElastiCache cluster and place it between the web application and RDS MySQL instances to reduce the load on the backend database tier.

Correct Answer: CE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 587
A user is using CloudFormation to launch an EC2 instance and then configure an application after the instance is launched.
The user wants the stack creation of ELB and AutoScaling to wait until the EC2 instance is launched and configured properly.
How can the user configure this?

A. The user can use the DependentCondition resource to hold the creation of the other dependent resources.

B. It is not possible that the stack creation will wait until one service is created and launched.
C. The user can use the HoldCondition resource to wait for the creation of the other dependent resources.
D. The user can use the WaitCondition resource to hold the creation of the other dependent resources.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
AWS CloudFormation is an application management tool that provides application modeling, deployment, configuration, management, and related activities.
AWS CloudFormation provides a WaitCondition resource that acts as a barrier and blocks the creation of other resources until a completion signal is received from an external source, such as a user application or management system.
Reference: http://aws.amazon.com/cloudformation/faqs

QUESTION 588
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 two.)

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 589
In Amazon Elastic Compute Cloud, you can specify storage volumes in addition to the root device volume when you create an AMI or when launching a new instance using_____.

A. block device mapping
B. object mapping
C. batch storage mapping
D. datacenter mapping

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
When creating an AMI or launching a new instance, you can assign more than one block storage device to it.
This device will be automatically set ready for you through an automated process known as block device mapping.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/block-device-mapping-concepts.html

QUESTION 590
An AWS account owner has setup multiple IAM users. One of these IAM users, named John, has CloudWatch access, but no access to EC2 services. John has setup an alarm action which stops EC2 instances when their CPU utilization is below the threshold limit.
When an EC2 instance's CPU Utilization rate drops below the threshold John has set, what will happen and why?

A. CloudWatch will stop the instance when the action is executed
B. Nothing will happen. John cannot set an alarm on EC2 since he does not have the permission.
C. Nothing will happen. John can setup the action, but it will not be executed because he does not have EC2 access through IAM policies.
D. Nothing will happen because it is not possible to stop the instance using the CloudWatch alarm

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Amazon CloudWatch alarms watch a single metric over a time period that the user specifies and performs one or more actions based on the value of the metric relative to a given threshold over a number of time periods. The user can setup an action which stops the instances when their CPU utilization is below a certain threshold for a certain period of time. The EC2 action can either terminate or stop the instance as part of the EC2 action. If the IAM user has read/write permissions for Amazon CloudWatch but not for Amazon EC2, he can still create an alarm.

However, the stop or terminate actions will not be performed on the Amazon EC2 instance.
Reference: http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/UsingAlarmActions.html

QUESTION 591
What is the maximum number of data points for an HTTP data request that a user can include in PutMetricRequest in theCloudWatch?

A. 30
B. 50
C. 10
D. 20

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The size of a PutMetricData request of CloudWatch is limited to 8KB for the HTTP GET requests and 40KB for the HTTPPOST requests. The user can include a maximum of 20 data points in one PutMetricData request.
Reference: http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/cloudwatch_concepts.html

QUESTION 592
An organization is planning to create a secure scalable application with AWS VPC and ELB. The organization has two instances already running and each instance has an ENI attached to it in addition to a primary network interface. The primary network interface and additional ENI both have an elastic IP attached to it.
If those instances are registered with ELB and the organization wants ELB to send data to a particular EIP of the instance, how can they achieve this?

A. The organization should ensure that the IP which is required to receive the ELB traffic is attached to a primary network interface.
B. It is not possible to attach an instance with two ENIs with ELB as it will give an IP conflict error.
C. The organization should ensure that the IP which is required to receive the ELB traffic is attached to an additional ENI.
D. It is not possible to send data to a particular IP as ELB will send to any one EIP.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web

Services (AWS) cloud. The user has complete control over the virtual networking environment.
Within this virtual private cloud, the user can launch AWS resources, such as an ELB, and EC2 instances. There are two ELBs available with VPC: internet facing and internal (private) ELB. For the internet facing ELB it is required that the ELB should be in a public subnet. When the user registers a multi-homed instance (an instance that has an Elastic Network Interface (ENI) attached) with a load balancer, the load balancer will route the traffic to the IP address of the primary network interface (eth0).
Reference: http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/gs-ec2VPC.html

QUESTION 593
A digital marketing company has multiple AWS accounts that belong to various teams. The creative team uses an Amazon S3 bucket in its AWS account to securely store images and media files that are used as content for the company's marketing campaigns. The creative team wants to share the S3 bucket with the strategy team so that the strategy team can view the objects.
A solutions architect has created an IAM role that is named strategy_reviewer in the Strategy account. The solutions architect also has set up a custom AWS Key Management Service (AWS KMS) key in the Creative account and has associated the key with the S3 bucket. However, when users from the Strategy account assume the IAM role and try to access objects in the S3 bucket, they receive an Access Denied error.
The solutions architect must ensure that users in the Strategy account can access the S3 bucket. The solution must provide these users with only the minimum permissions that they need.
Which combination of steps should the solutions architect take to meet these requirements? (Choose three.)

A. Create a bucket policy that includes read permissions for the S3 bucket. Set the principal of the bucket policy to the account ID of the Strategy account.
B. Update the strategy_reviewer IAM role to grant full permissions for the S3 bucket and to grant decrypt permissions for the custom KMS key.
C. Update the custom KMS key policy in the Creative account to grant decrypt permissions to the strategy_reviewer IAM role.
D. Create a bucket policy that includes read permissions for the S3 bucket. Set the principal of the bucket policy to an anonymous user.
E. Update the custom KMS key policy in the Creative account to grant encrypt permissions to the strategy_reviewer IAM role.
F. Update the strategy_reviewer IAM role to grant read permissions for the S3 bucket and to grant decrypt permissions for the custom KMS key.

Correct Answer: BDF
Section: (none)
Explanation

Explanation/Reference:

QUESTION 594
An advisory firm is creating a secure data analytics solution for its regulated financial services users. Users will upload their raw data to an Amazon S3 bucket, where they have PutObject permissions only. Data will be analyzed by applications running on an Amazon EMR cluster launched in a VPC. The firm requires that the environment be isolated from the internet.
All data at rest must be encrypted using keys controlled by the firm.
Which combination of actions should the Solutions Architect take to meet the user's security requirements? (Choose two.)

A. Launch the Amazon EMR cluster in a private subnet configured to use an AWS KMS CMK for at-rest encryption.

Configure a gateway VPC endpoint for Amazon S3 and an interface VPC endpoint for AWS KMS.

B. Launch the Amazon EMR cluster in a private subnet configured to use an AWS KMS CMK for at-rest encryption. Configure a gateway VPC endpoint for Amazon S3 and a NAT gateway to access AWS KMS.

C. Launch the Amazon EMR cluster in a private subnet configured to use an AWS CloudHSM appliance for at-rest encryption. Configure a gateway VPC endpoint for Amazon S3 and an interface VPC endpoint for CloudHSM.

D. Configure the S3 endpoint policies to permit access to the necessary data buckets only.

E. Configure the S3 bucket policies to permit access using an aws:sourceVpce condition to match the S3 endpoint ID.

Correct Answer: AE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 595
A company has a serverless multi-tenant content management system on AWS. The architecture contains a web-based front end that interacts with an Amazon API Gateway API that uses a custom AWS Lambda authorizer. The authorizer authenticates a user to its tenant ID and encodes the information in a JSON Web Token (JWT) token. After authentication, each API call through API Gateway targets a Lambda function that interacts with a single Amazon DynamoDB table to fulfill requests.
To comply with security standards, the company needs a stronger isolation between tenants. The company will have hundreds of customers within the first year.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create a DynamoDB table for each tenant by using the tenant ID in the table name. Create a service that uses the JWT token to retrieve the appropriate Lambda execution role that is tenant-specific. Attach IAM policies to the execution role to allow access only to the DynamoDB table for the tenant.

B. Add tenant ID information to the partition key of the DynamoDB table. Create a service that uses the JWT token to retrieve the appropriate Lambda execution role that is tenant-specific. Attach IAM policies to the execution role to allow access to items in the table only when the key matches the tenant ID.

C. Create a separate AWS account for each tenant of the application. Use dedicated infrastructure for each tenant. Ensure that no cross-account network connectivity exists.

D. Add tenant ID as a sort key in every DynamoDB table. Add logic to each Lambda function to use the tenant ID that comes from the JWT token as the sort key in every operation on the DynamoDB table.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://github.com/aws-samples/aws-saas-factory-dynamic-policy-generation

QUESTION 596
A retail company is running an application that stores invoice files in an Amazon S3 bucket and metadata about the files in an Amazon DynamoDB table. The application software runs in both us-east-1 and eu-west-1. The S3 bucket and DynamoDB table are in us-east-1. The company wants to protect itself from data corruption and loss of connectivity to either Region.
Which option meets these requirements?

A. Create a DynamoDB global table to replicate data between us-east-1 and eu-west-1. Enable continuous backup on the DynamoDB table in us-east-1. Enable versioning on the S3 bucket.
B. Create an AWS Lambda function triggered by Amazon CloudWatch Events to make regular backups of the DynamoDB table. Set up S3 cross-region replication from us-east-1 to eu-west-1. Set up MFA delete on the S3 bucket in us-east-1.
C. Create a DynamoDB global table to replicate data between us-east-1 and eu-west-1. Enable versioning on the S3 bucket. Implement strict ACLs on the S3 bucket.
D. Create a DynamoDB global table to replicate data between us-east-1 and eu-west-1. Enable continuous backup on the DynamoDB table in us-east-1. Set up S3 cross-region replication from us-east-1 to eu-west-1.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 597
A company wants to migrate its corporate data center from on premises to the AWS Cloud. The data center includes physical servers and VMs that use VMware and Hyper-V. An administrator needs to select the correct services to collect data for the initial migration discovery process. The data format should be supported by AWS Migration Hub. The company also needs the ability to generate reports from the data.
Which solution meets these requirements?

A. Use the AWS Agentless Discovery Connector for data collection on physical servers and all VMs. Store the collected data in Amazon S3. Query the data with S3 Select. Generate reports by using Kibana hosted on Amazon EC2.
B. Use the AWS Application Discovery Service agent for data collection on physical servers and all VMs. Store the collected data in Amazon Elastic File System (Amazon EFS). Query the data and generate reports with Amazon Athena.
C. Use the AWS Application Discovery Service agent for data collection on physical servers and Hyper-V. Use the AWS Agentless Discovery Connector for data collection on VMware. Store the collected data in Amazon S3. Query the data with Amazon Athena. Generate reports by using Amazon QuickSight.
D. Use the AWS Systems Manager agent for data collection on physical servers. Use the AWS Agentless Discovery Connector for data collection on all VMs. Store, query, and generate reports from the collected data by using Amazon Redshift.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 598
A company needs to architect a hybrid DNS solution. This solution will use an Amazon Route 53 private hosted zone for the domain cloud.example.com for the resources stored within VPCs.
The company has the following DNS resolution requirements:
On-premises systems should be able to resolve and connect to cloud.example.com.
All VPCs should be able to resolve cloud.example.com.
There is already an AWS Direct Connect connection between the on-premises corporate network and AWS Transit Gateway.
Which architecture should the company use to meet these requirements with the HIGHEST performance?

A. Associate the private hosted zone to all the VPCs. Create a Route 53 inbound resolver in the shared services VPC. Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver.
B. Associate the private hosted zone to all the VPCs. Deploy an Amazon EC2 conditional forwarder in the shared services VPC. Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the conditional forwarder.
C. Associate the private hosted zone to the shared services VPCreate a Route 53 outbound resolver in the shared services VPAttach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the outbound resolver.
D. Associate the private hosted zone to the shared services VPC. Create a Route 53 inbound resolver in the shared services VPC. Attach the shared services VPC to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/blogs/networking-and-content-delivery/centralized-dns-management-of-hybrid-cloudwith-amazon-route-53-and-aws-transit-gateway/

QUESTION 599
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 onpremises 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 onpremises 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 onpremises 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: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 600
An AWS customer has a web application that runs on premises. The web application fetches data from a third-party API that is behind a firewall. The third party accepts only one public CIDR block in each client's allow list.
The customer wants to migrate their web application to the AWS Cloud. The application will be hosted on a set of Amazon EC2 instances behind an Application Load Balancer (ALB) in a VPC. The ALB is located in public subnets. The EC2 instances are located in private subnets. NAT gateways provide internet access to the private subnets.
How should a solutions architect ensure that the web application can continue to call the third-party API after the migration?

A. Associate a block of customer-owned public IP addresses to the VPC. Enable public IP addressing for public subnets in the VPC.

B. Register a block of customer-owned public IP addresses in the AWS account. Create Elastic IP addresses from the address block and assign them to the NAT gateways in the VPC.

C. Create Elastic IP addresses from the block of customer-owned IP addresses. Assign the static Elastic IP addresses to the ALB.

D. Register a block of customer-owned public IP addresses in the AWS account. Set up AWS Global Accelerator to use Elastic IP addresses from the address block. Set the ALB as the accelerator endpoint.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 601
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:
Explanation:
Effect Deny

QUESTION 602
A company is manually deploying its application to production and wants to move to a more mature deployment pattern. The company has asked a solutions architect to design a solution that leverages its current Chef tools and knowledge. The application must be deployed to a staging environment for testing and verification before being deployed to production. Any new deployment must be rolled back in 5 minutes if errors are discovered after a deployment.
Which AWS service and deployment pattern should the solutions architect use to meet these requirements?

A. Use AWS Elastic Beanstalk and deploy the application using a rolling update deployment strategy.
B. Use AWS CodePipeline and deploy the application using a rolling update deployment strategy.
C. Use AWS CodeBuild and deploy the application using a canary deployment strategy.
D. Use AWS OpsWorks and deploy the application using a blue/green deployment strategy.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 603
A company runs a public-facing application that uses a Java-based web service via a RESTful API. It is hosted on Apache Tomcat on a single server in a data center that runs consistently at 30% CPU utilization. Use of the API is expected to increase by 10 times with a new product launch. The business wants to migrate the application to AWS with no disruption, and needs it to scale to meet demand.
The company has already decided to use Amazon Route 53 and CNAME records to redirect traffic. How can these requirements be met with the LEAST amount of effort?

A. Use AWS Elastic Beanstalk to deploy the Java web service and enable Auto Scaling. Then switch the application to use the new web service.
B. Lift and shift the Apache server to the cloud using AWS SMS. Then switch the application to direct web service traffic to the new instance.
C. Create a Docker image and migrate the image to Amazon ECS. Then change the application code to direct web service queries to the ECS container.
D. Modify the application to call the web service via Amazon API Gateway. Then create a new AWS Lambda Java function to run the Java web service code. After testing, change API Gateway to use the Lambda function.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 604
A company wants to refactor its retail ordering web application that currently has a load-balanced Amazon EC2 instance fleet for web hosting, database API services, and business logic. The company needs to create a decoupled, scalable architecture with a mechanism for retaining failed orders while also minimizing operational costs.
Which solution will meet these requirements?

A. Use Amazon S3 for web hosting with Amazon API Gateway for database API services. Use Amazon Simple Queue Service (Amazon SQS) for order queuing. Use Amazon Elastic Container Service (Amazon ECS) for business logic with Amazon SQS long polling for retaining failed orders.
B. Use AWS Elastic Beanstalk for web hosting with Amazon API Gateway for database API services. Use Amazon MQ for order queuing. Use AWS Step Functions for business logic with Amazon S3 Glacier Deep Archive for retaining failed orders.
C. Use Amazon S3 for web hosting with AWS AppSync for database API services. Use Amazon Simple Queue Service (Amazon SQS) for order queuing. Use AWS Lambda for business logic with an Amazon SQS dead-letter queue for retaining failed orders.
D. Use Amazon Lightsail for web hosting with AWS AppSync for database API services. Use Amazon Simple Email Service (Amazon SES) for order queuing. Use Amazon Elastic Kubernetes Service (Amazon EKS) for business logic with Amazon Elasticsearch Service (Amazon ES) for retaining failed orders.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 605
A company needs to cost-effectively persist small data records (up to 1 KiB) for up to 30 days. The data is read rarely. When reading the data, a 5-minute delay is acceptable.
Which of the following solutions achieve this goal? (Choose two.)

A. Use Amazon S3 to collect multiple records in one S3 object. Use a lifecycle configuration to move data to Amazon Glacier immediately after write. Use expedited retrievals when reading the data.
B. Write the records to Amazon Kinesis Data Firehose and configure Kinesis Data Firehose to deliver the data to Amazon S3 after 5 minutes. Set an expiration action at 30 days on the S3 bucket.
C. Use an AWS Lambda function invoked via Amazon API Gateway to collect data for 5 minutes. Write data to Amazon S3 just before the Lambda execution stops.
D. Write the records to Amazon DynamoDB configured with a Time To Live (TTL) of 30 days. Read data using the GetItem or BatchGetItem call.
E. Write the records to an Amazon ElastiCache for Redis. Configure the Redis append-only file (AOF) persistence logs to write to Amazon S3. Recover from the log if the ElastiCache instance has failed.

Correct Answer: AB
Section: (none)
Explanation

Explanation/Reference:

QUESTION 606
Can you configure multiple Load Balancers with a single Auto Scaling group?

A. No
B. Yes, you can but only if it is configured with Amazon Redshift.
C. Yes, you can provide the ELB is configured with Amazon AppStream.
D. Yes

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Yes, you can configure more than one load balancer with an autoscaling group. Auto Scaling integrates with Elastic Load Balancing to enable you to attach one or more load balancers to an existing Auto Scaling group. After you attach the load balancer, it automatically registers the instances in the group and distributes incoming traffic across the instances.
Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AS_Concepts.html

QUESTION 607
A Solutions Architect needs to design a highly available application that will allow authenticated users to stay connected to the application even when there are underlying failures.
Which solution will meet these requirements?

A. Deploy the application on Amazon EC2 instances. Use Amazon Route 53 to forward requests to the EC2 instances. Use Amazon DynamoDB to save the authenticated connection details.
B. Deploy the application on Amazon EC2 instances in an Auto Scaling group. Use an internet-facing Application Load Balancer to handle requests. Use Amazon DynamoDB to save the authenticated connection details.
C. Deploy the application on Amazon EC2 instances in an Auto Scaling group. Use an internet-facing Application Load Balancer on the front end. Use EC2 instances to save the authenticated connection details.
D. Deploy the application on Amazon EC2 instances in an Auto Scaling group. Use an internet-facing Application Load Balancer on the front end. Use EC2

instances hosting a MySQL database to save the authenticated connection details.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 608
A startup company recently migrated a large ecommerce website to AWS. The website has experienced a 70% increase in sales. Software engineers are using a private GitHub repository to manage code. The DevOps team is using Jenkins for builds and unit testing. The engineers need to receive notifications for bad builds and zero downtime during deployments.
The engineers also need to ensure any changes to production are seamless for users and can be rolled back in the event of a major issue.
The software engineers have decided to use AWS CodePipeline to manage their build and deployment process.
Which solution will meet these requirements?

A. Use GitHub websockets to trigger the CodePipeline pipeline. Use the Jenkins plugin for AWS CodeBuild to conduct unit testing. Send alerts to an Amazon SNS topic for any bad builds. Deploy in an inplace, all-at-once deployment configuration using AWS CodeDeploy.
B. Use GitHub webhooks to trigger the CodePipeline pipeline. Use the Jenkins plugin for AWS CodeBuild to conduct unit testing. Send alerts to an Amazon SNS topic for any bad builds. Deploy in a blue/ green deployment using AWS CodeDeploy.
C. Use GitHub websockets to trigger the CodePipeline pipeline. Use AWS X-Ray for unit testing and static code analysis. Send alerts to an Amazon SNS topic for any bad builds. Deploy in a blue/green deployment using AWS CodeDeploy.
D. Use GitHub webhooks to trigger the CodePipeline pipeline. Use AWS X-Ray for unit testing and static code analysis. Send alerts to an Amazon SNS topic for any bad builds. Deploy in an in-place, all-atonce deployment configuration using AWS CodeDeploy.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 609
Identify a benefit of using Auto Scaling for your application.

A. Your application gains better fault tolerance.
B. Your application optimizes only logistics and operations.
C. Your application receives latency requirements in every region.

D. You acquire clarity on prototypes in your application.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
When you use Auto Scaling, your applications gain better fault tolerance. Auto Scaling can detect when an instance is unhealthy, terminate it, and launch an instance to replace it. You can also configure Auto Scaling to use multiple Availability Zones. If one Availability Zone becomes unavailable, Auto Scaling can launch instances in another one to compensate.
Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/how-as-works.html

QUESTION 610
A three-tier web application runs on Amazon EC2 instances. Cron daemons are used to trigger scripts that collect the web server, application, and database logs and send them to a centralized location every hour. Occasionally, scaling events or unplanned outages have caused the instances to stop before the latest logs were collected, and the log files were lost.
Which of the following options is the MOST reliable way of collecting and preserving the log files?

A. Update the cron jobs to run every 5 minutes instead of every hour to reduce the possibility of log messages being lost in an outage.
B. Use Amazon CloudWatch Events to trigger Amazon Systems Manager Run Command to invoke the log collection scripts more frequently to reduce the possibility of log messages being lost in an outage.
C. Use the Amazon CloudWatch Logs agent to stream log messages directly to CloudWatch Logs. Configure the agent with a batch count of 1 to reduce the possibility of log messages being lost in an outage.
D. Use Amazon CloudWatch Events to trigger AWS Lambda to SSH into each running instance and invoke the log collection scripts more frequently to reduce the possibility of log messages being lost in an outage.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/AgentReference.html

QUESTION 611
A company CFO recently analyzed the company's AWS monthly bill and identified an opportunity to reduce the cost for AWS Elastic Beanstalk environments in use. The CFO has asked a Solutions Architect to design a highly available solution that will spin up an Elastic Beanstalk environment in the morning and terminate it at the end of the day.
The solution should be designed with minimal operational overhead and to minimize costs. It should also be able to handle the increased use of Elastic

Beanstalk environments among different teams, and must provide a one-stop scheduler solution for all teams to keep the operational costs low.
What design will meet these requirements?

A. Set up a Linux EC2 Micro instance. Configure an IAM role to allow the start and stop of the Elastic Beanstalk environment and attach it to the instance. Create scripts on the instance to start and stop the Elastic Beanstalk environment. Configure cron jobs on the instance to execute the scripts.
B. Develop AWS Lambda functions to start and stop the Elastic Beanstalk environment. Configure a Lambda execution role granting Elastic Beanstalk environment start/stop permissions, and assign the role to the Lambda functions. Configure cron expression Amazon CloudWatch Events rules to trigger the Lambda functions.
C. Develop an AWS Step Functions state machine with “wait” as its type to control the start and stop time. Use the activity task to start and stop the Elastic Beanstalk environment. Create a role for Step Functions to allow it to start and stop the Elastic Beanstalk environment. Invoke Step Functions daily.
D. Configure a time-based Auto Scaling group. In the morning, have the Auto Scaling group scale up an Amazon EC2 instance and put the Elastic Beanstalk environment start command in the EC2 instance user data. At the end of the day, scale down the instance number to 0 to terminate the EC2 instance.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://aws.amazon.com/premiumsupport/knowledge-center/schedule-elastic-beanstalk-stop-restart/

QUESTION 612
A company is planning to migrate its business-critical applications from an on-premises data center to AWS. The company has an on-premises installation of a Microsoft SQL Server Always On cluster. The company wants to migrate to an AWS managed database service. A solutions architect must design a heterogeneous database migration on AWS.
Which solution will meet these requirements?

A. Migrate the SQL Server databases to Amazon RDS for MySQL by using backup and restore utilities.
B. Use an AWS Snowball Edge Storage Optimized device to transfer data to Amazon S3. Set up Amazon RDS for MySQL.
Use S3 integration with SQL Server features, such as BULK INSERT.
C. Use the AWS Schema Conversion Tool to translate the database schema to Amazon RDS for MeSQL. Then use AWS Database Migration Service (AWS DMS) to migrate the data from on-premises databases to Amazon RDS.
D. Use AWS DataSync to migrate data over the network between on-premises storage and Amazon S3. Set up Amazon RDS for MySQL. Use S3 integration with SQL Server features, such as BULK INSERT.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/dms/latest/sbs/dms-sbs-welcome.html

QUESTION 613
A media company has a 30-TB repository of digital news videos. These videos are stored on tape in an on-premises tape library and referenced by a Media Asset Management (MAM) system. The company wants to enrich the metadata for these videos in an automated fashion and put them into a searchable catalog by using a MAM feature. The company must be able to search based on information in the video, such as objects, scenery items, or people's faces. A catalog is available that contains faces of people who have appeared in the videos that include an image of each person. The company would like to migrate these videos to AWS.
The company has a high-speed AWS Direct Connect connection with AWS and would like to move the MAM solution video content directly from its current file system.
How can these requirements be met by using the LEAST amount of ongoing management overhead and causing MINIMAL disruption to the existing system?

A. Set up an AWS Storage Gateway, file gateway appliance on-premises. Use the MAM solution to extract the videos from the current archive and push them into the file gateway. Use the catalog of faces to build a collection in Amazon Rekognition.
Build an AWS Lambda function that invokes the Rekognition Javascript SDK to have Rekognition pull the video from the Amazon S3 files backing the file gateway, retrieve the required metadata, and push the metadata into the MAM solution.
B. Set up an AWS Storage Gateway, tape gateway appliance on-premises. Use the MAM solution to extract the videos from the current archive and push them into the tape gateway. Use the catalog of faces to build a collection in Amazon Rekognition. Build an AWS Lambda function that invokes the Rekognition Javascript SDK to have Amazon Rekognition process the video in the tape gateway, retrieve the required metadata, and push the metadata into the MAM solution.
C. Configure a video ingestion stream by using Amazon Kinesis Video Streams. Use the catalog of faces to build a collection in Amazon Rekognition. Stream the videos from the MAM solution into Kinesis Video Streams. Configure Amazon Rekognition to process the streamed videos. Then, use a stream consumer to retrieve the required metadata, and push the metadata into the MAM solution. Configure the stream to store the videos in Amazon S3.
D. Set up an Amazon EC2 instance that runs the OpenCV libraries. Copy the videos, images, and face catalog from the onpremises library into an Amazon EBS volume mounted on this EC2 instance. Process the videos to retrieve the required metadata, and push the metadata into the MAM solution, while also copying the video files to an Amazon S3 bucket.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 614
A company has a 24 TB MySQL database in its on-premises data center that grows at the rate of 10 GB per day. The data center is connected to the company's AWS infrastructure with a 50 Mbps VPN connection.

The company is migrating the application and workload to AWS. The application code is already installed and tested on Amazon EC2. The company now needs to migrate the database and wants to go live on AWS within 3 weeks.
Which of the following approaches meets the schedule with LEAST downtime?

A. 1. Use the VM Import/Export service to import a snapshot of the on-premises database into AWS.
   2. Launch a new EC2 instance from the snapshot.
   3. Set up ongoing database replication from on premises to the EC2 database over the VPN.
   4. Change the DNS entry to point to the EC2 database.
   5. Stop the replication.

B. 1. Launch an AWS DMS instance.
   2. Launch an Amazon RDS Aurora MySQL DB instance.
   3. Configure the AWS DMS instance with on-premises and Amazon RDS database information.
   4. Start the replication task within AWS DMS over the VPN.
   5. Change the DNS entry to point to the Amazon RDS MySQL database.
   6. Stop the replication.

C. 1. Create a database export locally using database-native tools.
   2. Import that into AWS using AWS Snowball.
   3. Launch an Amazon RDS Aurora DB instance.
   4. Load the data in the RDS Aurora DB instance from the export.
   5. Set up database replication from the on-premises database to the RDS Aurora DB instance over the VPN.
   6. Change the DNS entry to point to the RDS Aurora DB instance.
   7. Stop the replication.

D. 1. Take the on-premises application offline.
   2. Create a database export locally using database-native tools.
   3. Import that into AWS using AWS Snowball.
   4. Launch an Amazon RDS Aurora DB instance.
   5. Load the data in the RDS Aurora DB instance from the export.
   6. Change the DNS entry to point to the Amazon RDS Aurora DB instance.
   7. Put the Amazon EC2 hosted application online.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 615
If I write the below command, what does it do? ec2-run ami-e3a5408a -n 20 -g appserver

A. Start twenty instances as members of appserver group.
B. Creates 20 rules in the security group named appserver
C. Terminate twenty instances as members of appserver group.
D. Start 20 security groups

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 616
A user has configured an EC2 instance in the US-East-1a zone. The user has enabled detailed monitoring of the instance.
The user is trying to get the data from CloudWatch using a CLI.
Which of the below mentioned CloudWatch endpoint URLs should the user use?

A. monitoring.us-east-1a.amazonaws.com
B. cloudwatch.us-east-1a.amazonaws.com
C. monitoring.us-east-1.amazonaws.com
D. monitoring.us-east-1-a.amazonaws.com

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The CloudWatch resources are always region specific and they will have the end point as region specific. If the user is trying to access the metric in the US-East-1 region, the endpoint URL will be: monitoring.us-east-1.amazonaws.com
Reference: http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/regions_endpoints.html

QUESTION 617
A company is building an image service on the web that will allow users to upload and search random photos. At peak usage, up to 10,000 users worldwide will upload their images. The service will then overlay text on the uploaded images, which will then be published on the company website.
Which design should a solutions architect implement?

A. Store the uploaded images in Amazon Elastic File System (Amazon EFS). Send application log information about each image to Amazon CloudWatch Logs. Create a fleet of Amazon EC2 instances that use CloudWatch Logs to determine which images need to be processed. Place processed images in another

directory in Amazon EFS Enable Amazon CloudFront and configure the origin to be the one of the EC2 instances in the fleet.

B. Store the uploaded images in an Amazon S3 bucket and configure an S3 bucket event notification to send a message to Amazon Simple Notification Service (Amazon SNS). Create a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB) to pull messages from Amazon SNS to process the images and place them in Amazon Elastic File System (Amazon EFS). Use Amazon CloudWatch metrics for the SNS message volume to scale out EC2 instances. Enable Amazon CloudFront and configure the origin to be the ALB in front of the EC2 instances.

C. Store the uploaded images in an Amazon S3 bucket and configure an S3 bucket event notification to send a message to the Amazon Simple Queue Service (Amazon SQS) queue. Create a fleet of Amazon EC2 instances to pull messages from the SQS queue to process the images and place them in another S3 bucket Use Amazon CloudWatch metrics for queue depth to scale out EC2 instances. Enable Amazon CloudFront and configure the origin to be the S3 bucket that contains the processed images.

D. Store the uploaded images on a shared Amazon Elastic Block Store (Amazon EBS) volume mounted to a fleet of Amazon EC2 Spot instances. Create an Amazon DynamoDB table that contains information about each uploaded image and whether it has been processed. Use an Amazon EventBridge (Amazon CloudWatch Events) rule to scale out EC2 instances. Enable Amazon CloudFront and configure the origin to reference an Elastic Load Balancer in front of the fleet of EC2 instances.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/AmazonS3/latest/userguide/NotificationHowTo.html

QUESTION 618
An organization is setting up a backup and restore system in AWS of their in premise system. The organization needs High Availability(HA) and Disaster Recovery(DR) but is okay to have a longer recovery time to save costs.
Which of the below mentioned setup options helps achieve the objective of cost saving as well as DR in the most effective way?

A. Setup pre-configured servers and create AMIs. Use EIP and Route 53 to quickly switch over to AWS from in premise.

B. Setup the backup data on S3 and transfer data to S3 regularly using the storage gateway.

C. Setup a small instance with AutoScaling; in case of DR start diverting all the load to AWS from on premise.

D. Replicate on premise DB to EC2 at regular intervals and setup a scenario similar to the pilot light.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
AWS has many solutions for Disaster Recovery(DR) and High Availability(HA). When the organization wants to have HA and DR but are okay to have a longer

recovery time they should select the option backup and restore with S3. The data can be sent to S3 using either Direct Connect, Storage Gateway or over the internet.
The EC2 instance will pick the data from the S3 bucket when started and setup the environment. This process takes longer but is very cost effective due to the low pricing of S3. In all the other options, the EC2 instance might be running or there will be AMI storage costs. Thus, it will be a costlier option. In this scenario the organization should plan appropriate tools to take a backup, plan the retention policy for data and setup security of the data.
Reference: http://d36cz9buwru1tt.cloudfront.net/AWS_Disaster_Recovery.pdf

QUESTION 619
A new startup is running a serverless application using AWS Lambda as the primary source of compute. New versions of the application must be made available to a subset of users before deploying changes to all users. Developers should also have the ability to abort the deployment and have access to an easy rollback mechanism. A solutions architect decides to use AWS CodeDeploy to deploy changes when a new version is available.
Which CodeDeploy configuration should the solutions architect use?

A. A blue/green deployment
B. A linear deployment
C. A canary deployment
D. An all-at-once deployment

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/automating-updates-to-serverlessapps.html

QUESTION 620
A financial services company receives a regular data feed from its credit card servicing partner. Approximately 5,000 records are sent every 15 minutes in plaintext, delivered over HTTPS directly into an Amazon S3 bucket with server-side encryption.
This feed contains sensitive credit card primary account number (PAN) data. The company needs to automatically mask the PAN before sending the data to another S3 bucket for additional internal processing. The company also needs to remove and merge specific fields, and then transform the record into JSON format. Additionally, extra feeds are likely to be added in the future, so any design needs to be easily expandable.
Which solutions will meet these requirements?

A. Trigger an AWS Lambda function on file delivery that extracts each record and writes it to an Amazon SQS queue.
   Trigger another Lambda function when new messages arrive in the SQS queue to process the records, writing the results to a temporary location in Amazon S3. Trigger a final Lambda function once the SQS queue is empty to transform the records into JSON format and send the results to another S3 bucket for internal processing.
B. Trigger an AWS Lambda function on file delivery that extracts each record and writes it to an Amazon SQS queue.

Configure an AWS Fargate container application to automatically scale to a single instance when the SQS queue contains messages. Have the application process each record, and transform the record into JSON format. When the queue is empty, send the results to another S3 bucket for internal processing and scale down the AWS Fargate instance.

C. Create an AWS Glue crawler and custom classifier based on the data feed formats and build a table definition to match. Trigger an AWS Lambda function on file delivery to start an AWS Glue ETL job to transform the entire record according to the processing and transformation requirements. Define the output format as JSON. Once complete, have the ETL job send the results to another S3 bucket for internal processing.

D. Create an AWS Glue crawler and custom classifier based upon the data feed formats and build a table definition to match. Perform an Amazon Athena query on file delivery to start an Amazon EMR ETL job to transform the entire record according to the processing and transformation requirements. Define the output format as JSON. Once complete, send the results to another S3 bucket for internal processing and scale down the EMR cluster.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 621
A company is planning the migration of several lab environments used for software testing. An assortment of custom tooling is used to manage the test runs for each lab. The labs use immutable infrastructure for the software test runs, and the results are stored in a highly available SQL database cluster. Although completely rewriting the custom tooling is out of scope for the migration project, the company would like to optimize workloads during the migration.
Which application migration strategy meets this requirement?

A. Re-host
B. Re-platform
C. Re-factor/re-architect
D. Retire

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://aws.amazon.com/blogs/enterprise-strategy/6-strategies-for-migrating-applications-to-the-cloud/

QUESTION 622
For Amazon EC2 issues, while troubleshooting AWS CloudFormation, you need to view the cloud-init and cfn logs for more information. Identify a directory to which these logs are published.

A. /var/opt/log/ec2
B. /var/log/lastlog
C. /var/log/
D. /var/log/ec2

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
When you use AWS CloudFormation, you might encounter issues when you create, update, or delete AWS CloudFormation stacks.
For Amazon EC2 issues, view the cloud-init and cfn logs. These logs are published on the Amazon EC2 instance in the /var/log/ directory. These logs capture processes and command outputs while AWS CloudFormation is setting up your instance. For Windows, view the EC2Configure service and cfn logs in % ProgramFiles%\Amazon\EC2ConfigService and C:\cfn\log.
You can also configure your AWS CloudFormation template so that the logs are published to Amazon CloudWatch, which displays logs in the AWS Management Console so you don't have to connect to your Amazon EC2 instance.
Reference: http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/troubleshooting.html

QUESTION 623
Which of the following should be followed before connecting to Amazon Virtual Private Cloud (Amazon VPC) using AWS Direct Connect?

A. Provide a public Autonomous System Number (ASN) to identify your network on the Internet.
B. Create a virtual private gateway and attach it to your Virtual Private Cloud (VPC).
C. Allocate a private IP address to your network in the 122.x.x.x range.
D. Provide a public IP address for each Border Gateway Protocol (BGP) session.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
To connect to Amazon Virtual Private Cloud (Amazon VPC) by using AWS Direct Connect, you must first do the following:
Provide a private Autonomous System Number (ASN) to identify your network on the Internet. Amazon then 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.
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.html

QUESTION 624
A Development team is deploying new APIs as serverless applications within a company. The team is currently using the AWS Management Console to provision Amazon API Gateway, AWS Lambda, and Amazon DynamoDB resources. A Solutions Architect has been tasked with automating the future deployments of these serverless APIs.
How can this be accomplished?

A. Use AWS CloudFormation with a Lambda-backed custom resource to provision API Gateway. Use the AWS::DynamoDB::Table and AWS::Lambda::Function resources to create the Amazon DynamoDB table and Lambda functions. Write a script to automate the deployment of the CloudFormation template.
B. Use the AWS Serverless Application Model to define the resources. Upload a YAML template and application files to the code repository. Use AWS CodePipeline to connect to the code repository and to create an action to build using AWS CodeBuild. Use the AWS CloudFormation deployment provider in CodePipeline to deploy the solution.
C. Use AWS CloudFormation to define the serverless application. Implement versioning on the Lambda functions and create aliases to point to the versions. When deploying, configure weights to implement shifting traffic to the newest version, and gradually update the weights as traffic moves over.
D. Commit the application code to the AWS CodeCommit code repository. Use AWS CodePipeline and connect to the CodeCommit code repository. Use AWS CodeBuild to build and deploy the Lambda functions using AWS CodeDeploy.
Specify the deployment preference type in CodeDeploy to gradually shift traffic over to the new version.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/quickstart/architecture/serverless-cicd-for-enterprise/ https://awsquickstart.s3.amazonaws.com/quickstart-trek10-serverless-enterprise-cicd/doc/serverless-cicd-for-the-enterprise-on-theaws-cloud.pdf

QUESTION 625
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: C

Section: (none)
Explanation

Explanation/Reference:

QUESTION 626
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: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 627
A software company is using three AWS accounts for each of its 10 development teams. The company has developed an AWS CloudFormation standard VPC template that includes three NAT gateways. The template is added to each account for each team. The company is concerned that network costs will increase each time a new development team is added. A solutions architect must maintain the reliability of the company's solutions and minimize operational complexity.
What should the solutions architect do to reduce the network costs while meeting these requirements?

A. Create a single VPC with three NAT gateways in a shared services account. Configure each account VPC with a default route through a transit gateway to the NAT gateway in the shared services account VPC. Remove all NAT gateways from the standard VPC template.
B. Create a single VPC with three NAT gateways in a shared services account. Configure each account VPC with a default route through a VPC peering connection to the NAT gateway in the shared services account VPC. Remove all NAT gateways from the standard VPC template.
C. Remove two NAT gateways from the standard VPC template. Rely on the NAT gateway SLA to cover reliability for the remaining NAT gateway.

D. Create a single VPC with three NAT gateways in a shared services account. Configure a Site-to-Site VPN connection from each account to the shared services account. Remove all NAT gateways from the standard VPC template.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 628
What is the role of the PollForTask action when it is called by a task runner in AWS Data Pipeline?

A. It is used to retrieve the pipeline definition.
B. It is used to report the progress of the task runner to AWS Data Pipeline.
C. It is used to receive a task to perform from AWS Data Pipeline.
D. It is used to inform AWS Data Pipeline of the outcome when the task runner completes a task.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Task runners call PollForTask to receive a task to perform from AWS Data Pipeline. If tasks are ready in the work queue, PollForTask returns a response immediately. If no tasks are available in the queue, PollForTask uses long-polling and holds on to a poll connection for up to 90 seconds, during which time any newly scheduled tasks are handed to the task agent.
Your remote worker should not call PollForTask again on the same worker group until it receives a response, and this may take up to 90 seconds.
Reference: http://docs.aws.amazon.com/datapipeline/latest/APIReference/API_PollForTask.html

QUESTION 629
A company's solutions architect is analyzing costs of a multi-application environment. The environment is deployed across multiple Availability Zones in a single AWS Region. After a recent acquisition, the company manages two organizations in AWS Organizations. The company has created multiple service provider applications as AWS PrivateLink-powered VPC endpoint services in one organization. The company has created multiple service consumer applications in the other organization.
Data transfer charges are much higher than the company expected, and the solutions architect needs to reduce the costs.
The solutions architect must recommend guidelines for developers to follow when they deploy services. These guidelines must minimize data transfer charges for the whole environment.
Which guidelines meet these requirements? (Choose two.)

A. Use AWS Resource Access Manager to share the subnets that host the service provider applications with other accounts in the organization.
B. Place the service provider applications and the service consumer applications in AWS accounts in the same organization.
C. Turn off cross-zone load balancing for the Network Load Balancer in all service provider application deployments.
D. Ensure that service consumer compute resources use the Availability Zone-specific endpoint service by using the endpoint's local DNS name.
E. Create a Savings Plan that provides adequate coverage for the organization's planned inter-Availability Zone data transfer usage.

Correct Answer: AC
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/ram/latest/userguide/getting-started-sharing.html
https://docs.aws.amazon.com/elasticloadbalancing/latest/classic/enable-disable-crosszone-lb.html

QUESTION 630
A company is running an Apache Hadoop cluster on Amazon EC2 instances. The Hadoop cluster stores approximately 100 TB of data for weekly operational reports and allows occasional access for data scientists to retrieve data. The company needs to reduce the cost and operational complexity for storing and serving this data.
Which solution meets these requirements in the MOST cost-effective manner?

A. Move the Hadoop cluster from EC2 instances to Amazon EMR. Allow data access patterns to remain the same.
B. Write a script that resizes the EC2 instances to a smaller instance type during downtime and resizes the instances to a larger instance type before the reports are created.
C. Move the data to Amazon S3 and use Amazon Athena to query the data for reports. Allow the data scientists to access the data directly in Amazon S3.
D. Migrate the data to Amazon DynamoDB and modify the reports to fetch data from DynamoDB. Allow the data scientists to access the data directly in DynamoDB.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 631
A group of research institutions and hospitals are in a partnership to study 2 PBs of genomic data. The institute that owns the data stores it in an Amazon S3 bucket and updates it regularly. The institute would like to give all of the organizations in the partnership read access to the data. All members of the partnership are extremely cost-conscious, and the institute that owns the account with the S3 bucket is concerned about covering the costs for requests and data transfers

from Amazon S3.
Which solution allows for secure datasharing without causing the institute that owns the bucket to assume all the costs for S3 requests and data transfers?

A. Ensure that all organizations in the partnership have AWS accounts. In the account with the S3 bucket, create a crossaccount role for each account in the partnership that allows read access to the data. Have the organizations assume and use that read role when accessing the data.
B. Ensure that all organizations in the partnership have AWS accounts. Create a bucket policy on the bucket that owns the data. The policy should allow the accounts in the partnership read access to the bucket. Enable Requester Pays on the bucket. Have the organizations use their AWS credentials when accessing the data.
C. Ensure that all organizations in the partnership have AWS accounts. Configure buckets in each of the accounts with a bucket policy that allows the institute that owns the data the ability to write to the bucket. Periodically sync the data from the institute's account to the other organizations. Have the organizations use their AWS credentials when accessing the data using their accounts.
D. Ensure that all organizations in the partnership have AWS accounts. In the account with the S3 bucket, create a crossaccount role for each account in the partnership that allows read access to the data. Enable Requester Pays on the bucket.
Have the organizations assume and use that read role when accessing the data.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 632
A group of Amazon EC2 instances have been configured as a high performance computing (HPC) cluster. The instances are running in a placement group, and are able to communicate with each other at network speeds of up to 20 Gbps.
The cluster needs to communicate with a control EC2 instance outside of the placement group. The control instance has the same instance type and AMI as the other instances, and is configured with a public IP address.
How can the Solutions Architect improve the network speeds between the control instance and the instances in the placement group?

A. Terminate the control instance and relaunch it in the placement group.
B. Ensure that the instances are communicating using their private IP addresses.
C. Ensure that the control instance is using an Elastic Network Adapter.
D. Move the control instance inside the placement group.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

Explanation:
Reference: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html

QUESTION 633
True or False: In Amazon ElastiCache, you can use Cache Security Groups to configure the cache clusters that are part of a VPC.

A. FALSE
B. TRUE
C. True, this is applicable only to cache clusters that are running in an Amazon VPC environment.
D. True, but only when you configure the cache clusters using the Cache Security Groups from the console navigation pane.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Amazon ElastiCache cache security groups are only applicable to cache clusters that are not running in an Amazon Virtual Private Cloud environment (VPC). If you are running in an Amazon Virtual Private Cloud, Cache Security Groups is not available in the console navigation pane.
Reference: http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/CacheSecurityGroup.html

QUESTION 634
An online retail company hosts its stateful web-based application and MySQL database in an on-premises data center on a single server. The company wants to increase its customer base by conducting more marketing campaigns and promotions.
In preparation, the company wants to migrate its application and database to AWS to increase the reliability of its architecture.
Which solution should provide the HIGHEST level of reliability?

A. Migrate the database to an Amazon RDS MySQL Multi-AZ DB instance. Deploy the application in an Auto Scaling group on Amazon EC2 instances behind an Application Load Balancer Store sessions in Amazon Neptune.
B. Migrate the database to Amazon Aurora MySQL. Deploy the application in an Auto Scaling group on Amazon EC2 instances behind an Application Load Balancer. Store sessions in an Amazon ElastiCache for Redis replication group.
C. Migrate the database to Amazon DocumentDB (with MongoDB compatibility). Deploy the application in an Auto Scaling group on Amazon EC2 instances behind a Network Load Balancer. Store sessions in Amazon Kinesis Data Firehose.
D. Migrate the database to an Amazon RDS MariaDB Multi-AZ DB instance. Deploy the application in an Auto Scaling group on Amazon EC2 instances behind an Application Load Balancer. Store sessions in Amazon ElastiCache for Memcached.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/documentdb/latest/developerguide/docdb-migration.html

QUESTION 635
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:
Explanation:
One ELB cannot handle different SSL certificates but since we are using sticky sessions it must be handled at the ELB level.
SSL could be handled on the EC2 instances only with TCP configured ELB, ELB supports sticky sessions only in HTTP/HTTPS configurations.
The way the Elastic Load Balancer does session stickiness is on a HTTP/HTTPS listener is by utilizing an HTTP cookie. IfSSL traffic is not terminated on the Elastic Load Balancer and is terminated on the back-end instance, the Elastic LoadBalancer has no visibility into the HTTP headers and therefore can not set or read any of the HTTP headers being passedback and forth. Reference:http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/elb-sticky-sessions.html

QUESTION 636
IAM users do not have permission to create Temporary Security Credentials for federated users and roles by default. In contrast, IAM users can call _____ without the need of any special permissions

A. GetSessionName
B. GetFederationToken
C. GetSessionToken

D. GetFederationName

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Currently the STS API command GetSessionToken is available to every IAM user in your account without previous permission. In contrast, the GetFederationToken command is restricted and explicit permissions need to be granted so a user can issue calls to this particular Action.
Reference: http://docs.aws.amazon.com/STS/latest/UsingSTS/STSPermission.html

QUESTION 637
A company hosts an application on Amazon EC2 instance and needs to store files in Amazon S3. The files should never traverse the public internet, and only the application EC2 instances are granted access to a specific Amazon S3 bucket. A solutions architect has created a VPC endpoint for Amazon S3 and connected the endpoint to the application VPC.
Which additional steps should the solutions architect take to meet these requirements?

A. Assign an endpoint policy to the endpoint that restricts access to a specific S3 bucket. Attach a bucket policy to the S3 bucket that grants access to the VPC endpoint. Add the gateway prefix list to a NACL of the instances to limit access to the application EC2 instances only.
B. Attach a bucket policy to the S3 bucket that grants access to application EC2 instances only using the aws:SourceIp condition. Update the VPC route table so only the application EC2 instances can access the VPC endpoint.
C. Assign an endpoint policy to the VPC endpoint that restricts access to a specific S3 bucket. Attach a bucket policy to the S3 bucket that grants access to the VPC endpoint. Assign an IAM role to the application EC2 instances and only allow access to this role in the S3 bucket's policy.
D. Assign an endpoint policy to the VPC endpoint that restricts access to S3 in the current Region. Attach a bucket policy to the S3 bucket that grants access to the VPC private subnets only. Add the gateway prefix list to a NACL to limit access to the application EC2 instances only.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 638
Your supervisor has given you the task of creating an elastic network interface on each of your web servers that connect to a mid-tier network where an application server resides. He also wants this set up as a Dual-homed Instance on Distinct Subnets. Instead of routing network packets through the dual-homed instances, where should each dual-homed instance receive and process requests to fulfil his criteria?

A. On one of the web servers

B. On the front end
C. On the back end
D. Through a security group

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
You can place an elastic network interface on each of your web servers that connects to a mid- tier network where an application server resides. The application server can also be dual-homed to a backend network (subnet) where the database server resides. If it is set up like this, instead of routing network packets through the dual-homed instances, each dual-homed instance receives and processes requests on the front end and initiates a connection to the back end before finally sending requests to the servers on the back-end network.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html

QUESTION 639
A company is configuring connectivity to a multi-account AWS environment to support application workloads that serve users in a single geographic region. The workloads depend on a highly available, onpremises legacy system deployed across two locations. It is critical for the AWS workloads to maintain connectivity to the legacy system, and a minimum of 5 Gbps of bandwidth is required. All application workloads within AWS must have connectivity with one another.
Which solution will meet these requirements?

A. Configure multiple AWS Direct Connect (DX) 10 Gbps dedicated connections from a DX partner for each on–premises location. Create private virtual interfaces on each connection for each AWS account VPC. Associate the private virtual interface with a virtual private gateway attached to each VPC.
B. Configure multiple AWS Direct Connect (DX) 10 Gbps dedicated connections from two DX partners for each on-premises location. Create and attach a virtual private gateway for each AWS account VPC. Create a DX gateway in a central network account and associate it with the virtual private gateways. Create a public virtual interface on each DX connection and associate the interface with the DX gateway.
C. Configure multiple AWS Direct Connect (DX) 10 Gbps dedicated connections from two DX partners for each on-premises location. Create a transit gateway and a DX gateway in a central network account. Create a transit virtual interface for each DX interface and associate them with the DX gateway. Create a gateway association between the DX gateway and the transit gateway.
D. Configure multiple AWS Direct Connect (DX) 10 Gbps dedicated connections from a DX partner for each on-premises location. Create and attach a virtual private gateway for each AWS account VPC. Create a transit gateway in a central network account and associate it with the virtual private gateways. Create a transit virtual interface on each DX connection and attach the interface to the transit gateway.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 640
The Statement element, of an AWS IAM policy, contains an array of individual statements. Each individual statement is a(n) _____ block enclosed in braces { }.

A. XML
B. JavaScript
C. JSON
D. AJAX

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The Statement element, of an IAM policy, contains an array of individual statements. Each individual statement is a JSON block enclosed in braces { }.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_ElementDescriptions.html

QUESTION 641
A Solutions Architect is building a containerized .NET Core application that will run in AWS Fargate. The backend of the application requires Microsoft SQL Server with high availability. All tiers of the application must be highly available. The credentials used for the connection string to SQL Server should not be stored on disk within the .NET Core front-end containers.
Which strategies should the Solutions Architect use to meet these requirements?

A. Set up SQL Server to run in Fargate with Service Auto Scaling. Create an Amazon ECS task execution role that allows the Fargate task definition to get the secret value for the credentials to SQL Server running in Fargate. Specify the ARN of the secret in AWS Secrets Manager in the secrets section of the Fargate task definition so the sensitive data can be injected into the containers as environment variables on startup for reading into the application to construct the connection string. Set up the .NET Core service using Service Auto Scaling behind an Application Load Balancer in multiple Availability Zones.
B. Create a Multi-AZ deployment of SQL Server on Amazon RDS. Create a secret in AWS Secrets Manager for the credentials to the RDS database. Create an Amazon ECS task execution role that allows the Fargate task definition to get the secret value for the credentials to the RDS database in Secrets Manager. Specify the ARN of the secret in Secrets Manager in the secrets section of the Fargate task definition so the sensitive data can be injected into the containers as environment variables on startup for reading into the application to construct the connection string. Set up the .NET Core service in Fargate using Service Auto Scaling behind an Application Load Balancer in multiple Availability Zones.
C. Create an Auto Scaling group to run SQL Server on Amazon EC2. Create a secret in AWS Secrets Manager for the credentials to SQL Server running on EC2. Create an Amazon ECS task execution role that allows the Fargate task definition to get the secret value for the credentials to SQL Server on EC2. Specify the ARN of the secret in Secrets Manager in the secrets section of the Fargate task definition so the sensitive data can be injected into the containers as environment variables on startup for reading into the application to construct the connection string. Set up the .NET Core service using Service Auto Scaling behind an Application Load Balancer in multiple Availability Zones.
D. Create a Multi-AZ deployment of SQL Server on Amazon RDS. Create a secret in AWS Secrets Manager for the credentials to the RDS database. Create

non-persistent empty storage for the .NET Core containers in the Fargate task definition to store the sensitive information. Create an Amazon ECS task execution role that allows the Fargate task definition to get the secret value for the credentials to the RDS database in Secrets Manager. Specify the ARN of the secret in Secrets Manager in the secrets section of the Fargate task definition so the sensitive data can be written to the nonpersistent empty storage on startup for reading into the application to construct the connection string. Set up the .NET Core service using Service Auto Scaling behind an Application Load Balancer in multiple Availability Zones.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 642
An organization is having an application which can start and stop an EC2 instance as per schedule. The organization needs the MAC address of the instance to be registered with its software. The instance is launched in EC2-CLASSIC.
How can the organization update the MAC registration every time an instance is booted?

A. The organization should write a boot strapping script which will get the MAC address from the instance metadata and use that script to register with the application.
B. The organization should provide a MAC address as a part of the user data. Thus, whenever the instance is booted the script assigns the fixed MAC address to that instance.
C. The instance MAC address never changes. Thus, it is not required to register the MAC address every time.
D. AWS never provides a MAC address to an instance; instead the instance ID is used for identifying the instance for any software registration.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
AWS provides an on demand, scalable infrastructure. AWS EC2 allows the user to launch On- Demand instances. AWS does not provide a fixed MAC address to the instances launched in EC2-CLASSIC. If the instance is launched as a part of EC2-VPC, it can have an ENI which can have a fixed MAC. However, with EC2-CLASSIC, every time the instance is started or stopped it will have a new MAC address. To get this MAC, the organization can run a script on boot which can fetch the instance metadata and get the MAC address from that instance metadata. Once the MAC is received, the organization can register that MAC with the software.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AESDG-chapter-instancedata.html

QUESTION 643
A company operates a group of imaging satellites. The satellites stream data to one of the company's ground stations where processing creates about 5 GB of

images per minute. This data is added to network-attached storage, where 2 PB of data are already stored.
The company runs a website that allows its customers to access and purchase the images over the Internet. This website is also running in the ground station.
Usage analysis shows that customers are most likely to access images that have been captured in the last 24 hours.
The company would like to migrate the image storage and distribution system to AWS to reduce costs and increase the number of customers that can be served.
Which AWS architecture and migration strategy will meet these requirements?

A. Use multiple AWS Snowball appliances to migrate the existing imagery to Amazon S3. Create a 1-Gb AWS Direct Connect connection from the ground station to AWS, and upload new data to Amazon S3 through the Direct Connect connection. Migrate the data distribution website to Amazon EC2 instances. By using Amazon S3 as an origin, have this website serve the data through Amazon CloudFront by creating signed URLs.
B. Create a 1-Gb Direct Connect connection from the ground station to AWS. Use the AWS Command Line Interface to copy the existing data and upload new data to Amazon S3 over the Direct Connect connection. Migrate the data distribution website to EC2 instances. By using Amazon S3 as an origin, have this website serve the data through CloudFront by creating signed URLs.
C. Use multiple Snowball appliances to migrate the existing images to Amazon S3. Upload new data by regularly using Snowball appliances to upload data from the network-attached storage. Migrate the data distribution website to EC2 instances. By using Amazon S3 as an origin, have this website serve the data through CloudFront by creating signed URLs.
D. Use multiple Snowball appliances to migrate the existing images to an Amazon EFS file system. Create a 1-Gb Direct Connect connection from the ground station to AWS, and upload new data by mounting the EFS file system over the Direct Connect connection. Migrate the data distribution website to EC2 instances. By using webservers in EC2 that mount the EFS file system as the origin, have this website serve the data through CloudFront by creating signed URLs.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 644
In AWS, which security aspects are the customer's responsibility? (Choose four.)

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: ACDF
Section: (none)

Explanation

Explanation/Reference:

QUESTION 645
Which of the following AWS services can be used to define alarms to trigger on a certain activity, such as activity success, failure, or delay in AWS Data Pipeline?

A. Amazon SES
B. Amazon CodeDeploy
C. Amazon SNS
D. Amazon SQS

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In AWS Data Pipeline, you can define Amazon SNS alarms to trigger on activities such as success, failure, or delay by creating an alarm object and referencing it in the onFail, onSuccess, or onLate slots of the activity object.
Reference:
https://aws.amazon.com/datapipeline/faqs/

QUESTION 646
A greeting card company recently advertised that customers could send cards to their favorite celebrities through the company's platform. Since the advertisement was published, the platform has received constant traffic from 10,000 unique users each second.
The platform runs on m5.xlarge Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an Auto Scaling group and use a custom AMI that is based on Amazon Linux. The platform uses a highly available Amazon Aurora MySQL DB cluster that uses primary and reader endpoints. The platform also uses an Amazon ElastiCache for Redis cluster that uses its cluster endpoint. The platform generates a new process for each customer and holds open database connections to MySQL for the duration of each customer's session. However, resource usage for the platform is low.
Many customers are reporting errors when they connect to the platform. Logs show that connections to the Aurora database are failing. Amazon CloudWatch metrics show that the CPU load is low across the platform and that connections to the platform are successful through the ALB.
Which solution will remediate the errors MOST cost-effectively?

A. Set up an Amazon CloudFront distribution. Set the ALB as the origin. Move all customer traffic to the CloudFront distribution endpoint.
B. Use Amazon RDS Proxy. Reconfigure the database connections to use the proxy.
C. Increase the number of reader nodes in the Aurora MySQL cluster.

D. Increase the number of nodes in the ElastiCache for Redis cluster.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 647
A company is running several large workloads on Amazon EC2 instances. Each EC2 instance has multiple Amazon Elastic Block Store (Amazon EBS) volumes attached to it. Once each day, an AWS Lambda function invokes the creation of EBS volume snapshots. These snapshots accumulate until an administrator manually purges them.
The company must maintain backups for a minimum of 30 days. A solutions architect needs to reduce the costs of this process.
Which solution meets these requirements MOST cost-effectively?

A. Search AWS Marketplace. Find a third-party solution to deploy to automatically manage the EBS volume backups.
B. Create a second Lambda function to move the EBS snapshots that are older than 30 days to Amazon S3 Glacier Deep Archive.
C. Set an Amazon S3 Lifecycle policy on the $3 bucket that contains the snapshots. Create a rule with an expiration action to delete EBS snapshots that are older than 30 days.
D. Migrate the backup functionality to Amazon Data Lifecycle Manager (Amazon DLM). Create a lifecycle policy for the daily backup of the EBS volumes. Set the retention period for the EBS snapshots to 30 days.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 648
A company that provides wireless services needs a solution to store and analyze log files about user activities. Currently, log files are delivered daily to Amazon Linux on an Amazon EC2 instance. A batch script is run once a day to aggregate data used for analysis by a third-party tool. The data pushed to the third-party tool is used to generate a visualization for end users. The batch script is cumbersome to maintain, and it takes several hours to deliver the ever-increasing data volumes to the third-party tool. The company wants to lower costs, and is open to considering a new tool that minimizes development effort and lowers administrative overhead. The company wants to build a more agile solution that can store and perform the analysis in near-real time, with minimal overhead. The solution needs to be cost effective and scalable to meet the company's end-user base growth.
Which solution meets the company's requirements?

A. Develop a Python script to capture the data from Amazon EC2 in real time and store the data in Amazon S3. Use a copy command to copy data from Amazon

S3 to Amazon Redshift. Connect a business intelligence tool running on Amazon EC2 to Amazon Redshift and create the visualizations.

B. Use an Amazon Kinesis agent running on an EC2 instance in an Auto Scaling group to collect and send the data to an Amazon Kinesis Data Firehose delivery stream. The Kinesis Data Firehose delivery stream will deliver the data directly to Amazon ES. Use Kibana to visualize the data.

C. Use an in-memory caching application running on an Amazon EBS-optimized EC2 instance to capture the log data in near real-time. Install an Amazon ES cluster on the same EC2 instance to store the log files as they are delivered to Amazon EC2 in near real-time. Install a Kibana plugin to create the visualizations.

D. Use an Amazon Kinesis agent running on an EC2 instance to collect and send the data to an Amazon Kinesis Data Firehose delivery stream. The Kinesis Data Firehose delivery stream will deliver the data to Amazon S3. Use an AWS Lambda function to deliver the data from Amazon S3 to Amazon ES. Use Kibana to visualize the data.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://docs.aws.amazon.com/firehose/latest/dev/writing-with-agents.html

QUESTION 649
A company wants to improve cost awareness for its Amazon EMR platform. The company has allocated budgets for each team's Amazon EMR usage. When a budgetary threshold is reached, a notification should be sent by email to the budget office's distribution list. Teams should be able to view their EMR cluster expenses to date. A solutions architect needs to create a solution that ensures the policy is proactively and centrally enforced in a multi-account environment. Which combination of steps should the solutions architect take to meet these requirements? (Choose two.)

A. Update the AWS CloudFormation template to include the AWS::Budgets::Budget::resource with the NotificationsWithSubscribers property.

B. Implement Amazon CloudWatch dashboards for Amazon EMR usage.

C. Create an EMR bootstrap action that runs at startup that calls the Cost Explorer API to set the budget on the cluster with the GetCostForecast and NotificationsWithSubscribers actions.

D. Create an AWS Service Catalog portfolio for each team. Add each team's Amazon EMR cluster as an AWS CloudFormation template to their Service Catalog portfolio as a Product.

E. Create an Amazon CloudWatch metric for billing. Create a custom alert when costs exceed the budgetary threshold.

Correct Answer: DE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 650
A company has an Amazon VPC that is divided into a public subnet and a private subnet. A web application runs in Amazon VPC, and each subnet has its own NACL. The public subnet has a CIDR of 10.0.0.0/24. An Application Load Balancer is deployed to the public subnet. The private subnet has a CIDR of 10.0.1.0/24. Amazon EC2 instances that run a web server on port 80 are launched into the private subnet.
Only network traffic that is required for the Application Load Balancer to access the web application can be allowed to travel between the public and private subnets. What collection of rules should be written to ensure that the private subnet's NACL meets the requirement? (Choose two.)

A. An inbound rule for port 80 from source 0.0.0.0/0.
B. An inbound rule for port 80 from source 10.0.0.0/24.
C. An outbound rule for port 80 to destination 0.0.0.0/0.
D. An outbound rule for port 80 to destination 10.0.0.0/24.
E. An outbound rule for ports 1024 through 65535 to destination 10.0.0.0/24.

Correct Answer: BC
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/vpc/latest/userguide/VPC_Scenario3.html

QUESTION 651
A Solutions Architect is working with a company that operates a standard three-tier web application in AWS. The web and application tiers run on Amazon EC2 and the database tier runs on Amazon RDS. The company is redesigning the web and application tiers to use Amazon API Gateway and AWS Lambda, and the company intends to deploy the new application within 6 months. The IT Manager has asked the Solutions Architect to reduce costs in the interim.
Which solution will be MOST cost effective while maintaining reliability?

A. Use Spot Instances for the web tier, On-Demand Instances for the application tier, and Reserved Instances for the database tier.
B. Use On-Demand Instances for the web and application tiers, and Reserved Instances for the database tier.
C. Use Spot Instances for the web and application tiers, and Reserved Instances for the database tier.
D. Use Reserved Instances for the web, application, and database tiers.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 652
Amazon Elastic File System (EFS) provides information about the space used for an object by using the space _ used attribute of the Network File System Version 4.1 (NFSv4.1). The attribute includes the object's current metered data size and not the metadata size. Which of the following utilities will you use to measure the amount of disk that is used of a file?

A. blkid utility
B. du utility
C. sfdisk utility
D. pydf utility

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Amazon EFS reports file system sizes and sizes of objects within a file system. Using the NFSv4.1 space _ used attribute for measuring the space used for an object, it reports only the object's current metered data size and not the metadata size.
There are two utilities available for measuring disk usage of a file, the du and stat utilities.
Reference:
https://docs.aws.amazon.com/efs/latest/ug/metered-sizes.html

QUESTION 653
A company is launching a web-based application in multiple regions around the world. The application consists of both static content stored in a private Amazon S3 bucket and dynamic content hosted in Amazon ECS containers content behind an Application Load Balancer (ALB). The company requires that the static and dynamic application content be accessible through Amazon CloudFront only. Which combination of steps should a solutions architect recommend to restrict direct content access to CloudFront? (Choose three.)

A. Create a web ACL in AWS WAF with a rule to validate the presence of a custom header and associate the web ACL with the ALB.
B. Create a web ACL in AWS WAF with a rule to validate the presence of a custom header and associate the web ACL with the CloudFront distribution.
C. Configure CloudFront to add a custom header to origin requests.
D. Configure the ALB to add a custom header to HTTP requests.
E. Update the S3 bucket ACL to allow access from the CloudFront distribution only.
F. Create a CloudFront Origin Access Identity (OAI) and add it to the CloudFront distribution. Update the S3 bucket policy to allow access to the OAI only.

Correct Answer: ADF
Section: (none)
Explanation

Explanation/Reference:

QUESTION 654
Which of the following cannot be used to manage Amazon ElastiCache and perform administrative tasks?

A. AWS software development kits (SDKs)
B. Amazon S3
C. ElastiCache command line interface (CLI)
D. AWS CloudWatch

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
CloudWatch is a monitoring tool and doesn't give users access to manage Amazon ElastiCache.
Reference: http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/WhatIs.Managing.html

QUESTION 655
A company is serving files to its customer through an SFTP server that is accessible over the Internet. The SFTP server is running on a single Amazon EC2 instance with an Elastic IP address attached. Customers connect to the SFTP server through its Elastic IP address and use SSH for authentication. The EC2 instance also has an attached security group that allows access from all customer IP addresses.
A solutions architect must implement a solution to improve availability, minimize the complexity of infrastructure management, and minimize the disruption to customers who access files. The solution must not change the way customers connect.
Which solution will meet these requirements?

A. Disassociate the Elastic IP address from the EC2 instance. Create an Amazon S3 bucket to be used for SFTP file hosting. Create an AWS Transfer Family server. Configure the Transfer Family server with a publicly accessible endpoint.
Associate the SFTP Elastic IP address with the new endpoint. Point the Transfer Family server to the S3 bucket. Sync all files from the SFTP server to the S3 bucket.
B. Disassociate the Elastic IP address from the EC2 instance. Create an Amazon S3 bucket to be used for SFTP file hosting. Create an AWS Transfer Family server. Configure the Transfer Family server with a VPC-hosted, Internet-facing endpoint. Associate the SFTP Elastic IP address with the new endpoint. Attach the security group with customer IP addresses to the new endpoint. Point the Transfer Family server to the S3 bucket. Sync all files from the SFTP server to the S3 bucket.
C. Disassociate the Elastic IP address from the EC2 instance. Create a new Amazon Elastic File System (Amazon EFS) file system to be used for SFTP file hosting. Create an AWS Fargate task definition to run an SFTP server. Specify the EFS file system as a mount in the task definition. Create a Fargate service by using the task definition, and place a Network Load Balancer (NLB) in front of the service. When configuring the service, attach the security group with

customer IP addresses to the tasks that run the SFTP server. Associate the Elastic IP address with the NLB. Sync all files from the SFTP server to the S3 bucket.

D. Disassociate the Elastic IP address from the EC2 instance. Create a multi-attach Amazon Elastic Block Store (Amazon EBS) volume to be used for SFTP file hosting. Create a Network Load Balancer (NLB) with the Elastic IP address attached.
Create an Auto Scaling group with EC2 instances that run an SFTP server. Define in the Auto Scaling group that instances that are launched should attach the new multi-attach EBS volume. Configure the Auto Scaling group to automatically add instances behind the NLB. Configure the Auto Scaling group to use the security group that allows customer IP addresses for the EC2 instances that the Auto Scaling group launches. Sync all files from the SFTP server to the new multi-attach EBS volume.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 656
ABC has three separate departments and each department has their own AWS accounts. The HR department has created a file sharing site where all the on roll employees' data is uploaded. The Admin department uploads data about the employee presence in the office to their DB hosted in the VPC. The Finance department needs to access data from the HR department to know the on roll employees to calculate the salary based on the number of days that an employee is present in the office.
How can ABC setup this scenario?

A. It is not possible to configure VPC peering since each department has a separate AWS account.
B. Setup VPC peering for the VPCs of Admin and Finance.
C. Setup VPC peering for the VPCs of Finance and HR as well as between the VPCs of Finance and Admin.
D. Setup VPC peering for the VPCs of Admin and HR

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. A VPC peering connection allows the user to route traffic between the peer VPCs using private IP addresses as if they are a part of the same network. This is helpful when one VPC from the same or different AWS account wants to connect with resources of the other VPC.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/PeeringGuide/peering-configurations-full-access.html#three-vpcs-fullaccess.

QUESTION 657

An organization is having a VPC for the HR department, and another VPC for the Admin department. The HR department requires access to all the instances running in the Admin VPC while the Admin department requires access to all the resources in the HR department.
How can the organization setup this scenario?

A. Setup VPC peering between the VPCs of Admin and HR.
B. Setup ACL with both VPCs which will allow traffic from the CIDR of the other VPC.
C. Setup the security group with each VPC which allows traffic from the CIDR of another VPC.
D. It is not possible to connect resources of one VPC from another VPC.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. A VPC peering connection allows the user to route traffic between the peer VPCs using private IP addresses as if they are a part of the same network. This is helpful when one VPC from the same or different AWS account wants to connect with resources of the other VPC.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpc-peering.html

QUESTION 658
A company wants to run a serverless application on AWS. The company plans to provision its application in Docker containers running in an Amazon ECS cluster. The application requires a MySQL database and the company plans to use Amazon RDS. The company has documents that need to be accessed frequently for the first 3 months, and rarely after that.
The document must be retained for 7 years.
What is the MOST cost-effective solution to meet these requirements?

A. Create an ECS cluster using On-Demand Instances. Provision the database and its read replicas in Amazon RDS using Spot Instances. Store the documents in an encrypted EBS volume, and create a cron job to delete the documents after 7 years.
B. Create an ECS cluster using a fleet of Spot Instances, with Spot Instance draining enabled. Provision the database and its read replicas in Amazon RDS using Reserved Instances. Store the documents in a secured Amazon S3 bucket with a lifecycle policy to move the documents that are older than 3 months to Amazon S3 Glacier, then delete the documents from Amazon S3 Glacier that are more than 7 years old.
C. Create an ECS cluster using On-Demand Instances. Provision the database and its read replicas in Amazon RDS using On-Demand Instances. Store the documents in Amazon EFS. Create a cron job to move the documents that are older than 3 months to Amazon S3 Glacier. Create an AWS Lambda function to delete the documents in S3 Glacier that are older than 7 years.
D. Create an ECS cluster using a fleet of Spot Instances with Spot Instance draining enabled. Provision the database and its read replicas in Amazon RDS using On-Demand Instances. Store the documents in a secured Amazon S3 bucket with a lifecycle policy to move the documents that are older than 3 months to Amazon S3 Glacier, then delete the documents in Amazon S3 Glacier after 7 years.

Correct Answer: B

Section: (none)
Explanation

Explanation/Reference:

QUESTION 659
A user is trying to send custom metrics to CloudWatch using the PutMetricData APIs. Which of the below mentioned points should the user needs to take care while sending the data to CloudWatch?

A. The size of a request is limited to 8KB for HTTP GET requests and 40KB for HTTP POST requests
B. The size of a request is limited to 16KB for HTTP GET requests and 80KB for HTTP POST requests
C. The size of a request is limited to 128KB for HTTP GET requests and 64KB for HTTP POST requests
D. The size of a request is limited to 40KB for HTTP GET requests and 8KB for HTTP POST requests

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
With AWS CloudWatch, the user can publish data points for a metric that share not only the same time stamp, but also the same namespace and dimensions. CloudWatch can accept multiple data points in the same PutMetricData call with the same time stamp. The only thing that the user needs to take care of is that the size of a PutMetricData request is limited to 8KB for HTTP GET requests and 40KB for HTTP POST requests.
Reference: http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/cloudwatch_concepts.html

QUESTION 660
A financial services company has an on-premises environment that ingests market data feeds from stock exchanges, transforms the data, and sends the data to an internal Apache Kafka cluster. Management wants to leverage AWS services to build a scalable and near-real-time solution with consistent network performance to provide stock market data to a web application.
Which steps should a solutions architect take to build the solution? (Choose three.)

A. Establish an AWS Direct Connect connection from the on-premises data center to AWS.
B. Create an Amazon EC2 Auto Scaling group to pull the messages from the on-premises Kafka cluster and use the Amazon Consumer Library to put the data into an Amazon Kinesis data stream.
C. Create an Amazon EC2 Auto Scaling group to pull the messages from the on-premises Kafka cluster and use the Amazon Kinesis Producer Library to put the data into a Kinesis data stream.
D. Create a WebSocket API in Amazon API Gateway, create an AWS Lambda function to process an Amazon Kinesis data stream, and use the @connections command to send callback messages to connected clients.

E. Create a GraphQL API in AWS AppSync, create an AWS Lambda function to process the Amazon Kinesis data stream, and use the @connections command to send callback messages to connected clients.
F. Establish a Site-to-Site VPN from the on-premises data center to AWS.

Correct Answer: ADE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 661
An organization has two Amazon EC2 instances:
The first is running an ordering application and an inventory application. The second is running a queuing system.
During certain times of the year, several thousand orders are placed per second. Some orders were lost when the queuing system was down. Also, the organization's inventory application has the incorrect quantity of products because some orders were processed twice.
What should be done to ensure that the applications can handle the increasing number of orders?

A. Put the ordering and inventory applications into their own AWS Lambda functions. Have the ordering application write the messages into an Amazon SQS FIFO queue.
B. Put the ordering and inventory applications into their own Amazon ECS containers, and create an Auto Scaling group for each application. Then, deploy the message queuing server in multiple Availability Zones.
C. Put the ordering and inventory applications into their own Amazon EC2 instances, and create an Auto Scaling group for each application. Use Amazon SQS standard queues for the incoming orders, and implement idempotency in the inventory application.
D. Put the ordering and inventory applications into their own Amazon EC2 instances. Write the incoming orders to an Amazon Kinesis data stream. Configure AWS Lambda to poll the stream and update the inventory application.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/standard-queues.html

QUESTION 662
While debugging a backend application for an IoT system that supports globally distributed devices, a Solutions Architect notices that stale data is occasionally being sent to user devices. Devices often share data, and stale data does not cause issues in most cases. However, device operations are disrupted when a device reads the stale data after an update.

The global system has multiple identical application stacks deployed in different AWS Regions. If a user device travels out of its home geographic region, it will always connect to the geographically closest AWS Region to write or read data. The same data is available in all supported AWS Regions using an Amazon DynamoDB global table.
What change should be made to avoid causing disruptions in device operations?

A. Update the backend to use strongly consistent reads. Update the devices to always write to and read from their home AWS Region.
B. Enable strong consistency globally on a DynamoDB global table. Update the backend to use strongly consistent reads.
C. Switch the backend data store to Amazon Aurora MySQL with cross-region replicas. Update the backend to always write to the master endpoint.
D. Select one AWS Region as a master and perform all writes in that AWS Region only. Update the backend to use strongly consistent reads.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 663
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:
Explanation:

Reference:
https://aws.amazon.com/blogs/aws/elastic-load-balancing-adds-support-for-proxy-protocol/

QUESTION 664
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 sharedstorage 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 upto- date standby. Compare this to a Single-AZ deployment: in case of a Single-AZ database failure, a user-initiated point-intime- 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 665
A mobile gaming application publishes data continuously to Amazon Kinesis Data Streams. An AWS Lambda function processes records from the data stream and writes to an Amazon DynamoDB table. The DynamoDB table has an auto scaling policy enabled with the target utilization set to 70%.
For several minutes at the start and end of each day, there is a spike in traffic that often exceeds five times the normal load.
The company notices the GetRecords.IteratorAgeMilliseconds metric of the Kinesis data stream temporarily spikes to over a minute for several minutes. The AWS Lambda function writes ProvisionedThroughputExceededException messages to Amazon CloudWatch Logs during these times, and some records are redirected to the dead letter queue. No exceptions are thrown by the Kinesis producer on the gaming application.
What change should the company make to resolve this issue?

A. Use Application Auto Scaling to set a scaling schedule to scale out write capacity on the DynamoDB table during predictable load spikes.
B. Use Amazon CloudWatch Events to monitor the dead letter queue and invoke a Lambda function to automatically retry failed records.
C. Reduce the DynamoDB table auto scaling policy's target utilization to 20% to more quickly respond to load spikes.
D. Increase the number of shards in the Kinesis data stream to increase throughput capacity.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 666
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.
E. Can I use AWS Direct Connect and a VPN Connection to the same VPC simultaneously?
Yes. However, only in fail-over scenarios. The Direct Connect path will always be preferred, when established, regardless of AS path prepending. Reference: https://aws.amazon.com/directconnect/faqs/

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Q. Can I use AWS Direct Connect and a VPN Connection to the same VPC simultaneously?
Yes. However, only in fail-over scenarios. The Direct Connect path will always be preferred, when established, regardless of AS path prepending. Reference: https://aws.amazon.com/directconnect/faqs/

QUESTION 667
A retail company processes point-of-sale data on application servers in its data center and writes outputs to an Amazon DynamoDB table. The data center is connected to the company's VPC with an AWS Direct Connect (DX) connection, and the application servers require a consistent network connection at speeds greater than 2 Gbps.
The company decides that the DynamoDB table needs to be highly available and fault tolerant. The company policy states that the data should be available across two regions.
What changes should the company make to meet these requirements?

A. Establish a second DX connection for redundancy. Use DynamoDB global tables to replicate data to a second Region. Modify the application to fail over to the second Region.
B. Use an AWS managed VPN as a backup to DX. Create an identical DynamoDB table in a second Region. Modify the application to replicate data to both Regions.
C. Establish a second DX connection for redundancy. Create an identical DynamoDB table in a second Region. Enable DynamoDB auto scaling to manage throughput capacity. Modify the application to write to the second Region.
D. Use AWS managed VPN as a backup to DX. Create an identical DynamoDB table in a second Region. Enable DynamoDB streams to capture changes to the table. Use AWS Lambda to replicate changes to the second Region.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 668
A user is trying to understand the detailed CloudWatch monitoring concept. Which of the below mentioned services does not provide detailed monitoring with CloudWatch?

A. AWS RDS
B. AWS ELB
C. AWS Route53
D. AWS EMR

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:

CloudWatch is used to monitor AWS as well as the custom services. It provides either basic or detailed monitoring for the supported AWS products. In basic monitoring, a service sends data points to CloudWatch every five minutes, while in detailed monitoring a service sends data points to CloudWatch every minute. Services, such as RDS, EC2, Auto Scaling, ELB, and Route 53 can provide the monitoring data every minute.
Reference: http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/supported_services.html

QUESTION 669
A company runs a legacy system on a single m4.2xlarge Amazon EC2 instance with Amazon EBS storage. The EC2 instance runs both the web server and a self-managed Oracle database. A snapshot is made of the EBS volume every 12 hours, and an AMI was created from the fully configured EC2 instance.
A recent event that terminated the EC2 instance led to several hours of downtime. The application was successfully launched from the AMI, but the age of the EBS snapshot and the repair of the database resulted in the loss of 8 hours of data. The system was also down for 4 hours while the Systems Operators manually performed these processes. What architectural changes will minimize downtime and reduce the chance of lost data?

A. Create an Amazon CloudWatch alarm to automatically recover the instance. Create a script that will check and repair the database upon reboot. Subscribe the Operations team to the Amazon SNS message generated by the CloudWatch alarm.
B. Run the application on m4.xlarge EC2 instances behind an Elastic Load Balancer/Application Load Balancer. Run the EC2 instances in an Auto Scaling group across multiple Availability Zones with a minimum instance count of two. Migrate the database to an Amazon RDS Oracle Multi-AZ DB instance.
C. Run the application on m4.2xlarge EC2 instances behind an Elastic Load Balancer/Application Load Balancer. Run the EC2 instances in an Auto Scaling group across multiple Availability Zones with a minimum instance count of one. Migrate the database to an Amazon RDS Oracle Multi-AZ DB instance.
D. Increase the web server instance count to two m4.xlarge instances and use Amazon Route 53 round-robin load balancing to spread the load. Enable Route 53 health checks on the web servers. Migrate the database to an Amazon RDS Oracle Multi-AZ DB instance.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html

QUESTION 670
Someone has recommended a new client to you and you know he is into online gaming and you are almost certain he will want to set up an online gaming site which will require a database service that provides fast and predictable performance with seamless scalability.
Which of the following AWS databases would be best suited to an online gaming site?

A. Amazon SimpleDB
B. Amazon DynamoDB
C. Amazon Redshift
D. Amazon ElastiCache

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Amazon DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. You can use Amazon DynamoDB to create a database table that can store and retrieve any amount of data, and serve any level of request traffic. Amazon DynamoDB automatically spreads the data and traffic for the table over a sufficient number of servers to handle the request capacity specified by the customer and the amount of data stored, while maintaining consistent and fast performance.
Reference: http://aws.amazon.com/documentation/dynamodb/

QUESTION 671
A company has a primary Amazon S3 bucket that receives thousands of objects every day. The company needs to replicate these objects into several other S3 buckets from various AWS accounts. A solutions architect is designing a new AWS Lambda function that is triggered when an object is created in the main bucket and replicates the object into the target buckets. The objects do not need to be replicated in real time. There is concern that this function may impact other critical Lambda functions due to Lambda's regional concurrency limit.
How can the solutions architect ensure this new Lambda function will not impact other critical Lambda functions?

A. Set the new Lambda function reserved concurrency limit to ensure the executions do not impact other critical Lambda functions. Monitor existing critical Lambda functions with Amazon CloudWatch alarms for the Throttles Lambda metric.
B. Increase the execution timeout of the new Lambda function to 5 minutes. Monitor existing critical Lambda functions with Amazon CloudWatch alarms for the Throttles Lambda metric.
C. Configure S3 event notifications to add events to an Amazon SQS queue in a separate account. Create the new Lambda function in the same account as the SQS queue and trigger the function when a message arrives in the queue.
D. Ensure the new Lambda function implements an exponential backoff algorithm. Monitor existing critical Lambda functions with Amazon CloudWatch alarms for the Throttles Lambda metric.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 672
How can you check the operational validity of your AWS CloudFormation template?

A. To check the operational validity, you need to attempt to create the stack.
B. There is no way to check the operational validity of your AWS CloudFormation template.

C. To check the operational validity, you need a sandbox or test area for AWS CloudFormation stacks.
D. To check the operational validity, you need to use the aws cloudformation validate-template command.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In AWS CloudFormation, to check the operational validity, you need to attempt to create the stack. There is no sandbox or test area for AWS CloudFormation stacks, so you are charged for the resources you create during testing.
Reference: http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-validate-template.html

QUESTION 673
A company has an application that runs on a fleet of Amazon EC2 instances and stores 70 GB of device data for each instance in Amazon S3. Recently, some of the S3 uploads have been failing. At the same time, the company is seeing an unexpected increase in storage data costs. The application code cannot be modified.
What is the MOST efficient way to upload the device data to Amazon S3 while managing storage costs?

A. Upload device data using a multipart upload. Use the AWS CLI to list incomplete parts to address the failed S3 uploads.
Enable the lifecycle policy for the incomplete multipart uploads on the S3 bucket to delete the old uploads and prevent new failed uploads from accumulating.
B. Upload device data using S3 Transfer Acceleration. Use the AWS Management Console to address the failed S3 uploads. Use the Multi-Object Delete operation nightly to delete the old uploads.
C. Upload device data using a multipart upload. Use the AWS Management Console to list incomplete parts to address the failed S3 uploads. Configure a lifecycle policy to archive continuously to Amazon S3 Glacier.
D. Upload device data using S3 Transfer Acceleration. Use the AWS Management Console to list incomplete parts to address the failed S3 uploads. Enable the lifecycle policy for the incomplete multipart uploads on the S3 bucket to delete the old uploads and prevent new failed uploads from accumulating.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/amazonglacier/latest/dev/uploading-an-archive.html

QUESTION 674
The _____ service is targeted at organizations with multiple users or systems that use AWS products such as Amazon EC2, Amazon SimpleDB, and the AWS Management Console.

A. Amazon RDS
B. AWS Integrity Management
C. AWS Identity and Access Management
D. Amazon EMR

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/documentation/iam/?nc1=h_ls

QUESTION 675
An Auto Scaling group is running at the desired capacity of 5 instances and receives a trigger from the Cloudwatch Alarm to increase the capacity by 1. The cool down period is 5 minutes. Cloudwatch sends another trigger after 2 minutes to decrease the desired capacity by 1.
What will be the count of instances at the end of 4 minutes?

A. 4
B. 5
C. 6
D. 7

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The cool down period is the time difference between the end of one scaling activity (can be start or terminate) and the start of another one (can be start or terminate). During the cool down period, Auto Scaling does not allow the desired capacity of the Auto Scaling group to be changed by any other CloudWatch alarm. Thus, in this case the trigger from the second alarm will have no effect.
Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AS_Concepts.html#healthcheck

QUESTION 676
A company runs a three-tier application in AWS. Users report that the application performance can vary greatly depending on the time of day and functionality being accessed.
The application includes the following components:
Eight t2.large front-end web servers that serve static content and proxy dynamic content from the application tier. Four t2.large application servers.

One db.m4.large Amazon RDS MySQL Multi-AZ DB instance.
Operations has determined that the web and application tiers are network constrained.
Which of the following is a cost effective way to improve application performance? (Choose two.)

A. Replace web and app tiers with t2.xlarge instances
B. Use AWS Auto Scaling and m4.large instances for the web and application tiers
C. Convert the MySQL RDS instance to a self-managed MySQL cluster on Amazon EC2
D. Create an Amazon CloudFront distribution to cache content
E. Increase the size of the Amazon RDS instance to db.m4.xlarge

Correct Answer: BD
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/ec2/instance-types/

QUESTION 677
A solutions architect is importing a VM from an on-premises environment by using the Amazon EC2 VM Import feature of AWS Import/Export. The solutions architect has created an AMI and has provisioned an Amazon EC2 instance that is based on that AMI. The EC2 instance runs inside a public subnet in a VPC and has a public IP address assigned.
The EC2 instance does not appear as a managed instance in the AWS Systems Manager console.
Which combination of steps should the solutions architect take to troubleshoot this issue? (Choose two.)

A. Verify that Systems Manager Agent is installed on the instance and is running
B. Verify that the instance is assigned an appropriate IAM role for Systems Manager.
C. Verify the existence of a VPC endpoint on the VPC.
D. Verify that the AWS Application Discovery Agent is configured.
E. Verify the correct configuration of service-linked roles for Systems Manager.

Correct Answer: BC
Section: (none)
Explanation

Explanation/Reference:

QUESTION 678

An enterprise company's data science team wants to provide a safe, cost-effective way to provide easy access to Amazon SageMaker. The data scientists have limited AWS knowledge and need to be able to launch a Jupyter notebook instance.
The notebook instance needs to have a preconfigured AWS KMS key to encrypt data at rest on the machine learning storage volume without exposing the complex setup requirements.
Which approach will allow the company to set up a self-service mechanism for the data scientists to launch Jupyter notebooks in its AWS accounts with the LEAST amount of operational overhead?

A. Create a serverless front end using a static Amazon S3 website to allow the data scientists to request a Jupyter notebook instance by filling out a form. Use Amazon API Gateway to receive requests from the S3 website and trigger a central AWS Lambda function to make an API call to Amazon SageMaker that will launch a notebook instance with a preconfigured KMS key for the data scientists. Then call back to the front-end website to display the URL to the notebook instance.
B. Create an AWS CloudFormation template to launch a Jupyter notebook instance using the AWS::SageMaker::NotebookInstance resource type with a preconfigured KMS key. Add a user-friendly name to the CloudFormation template. Display the URL to the notebook using the Outputs section. Distribute the CloudFormation template to the data scientists using a shared Amazon S3 bucket.
C. Create an AWS CloudFormation template to launch a Jupyter notebook instance using the AWS::SageMaker::NotebookInstance resource type with a preconfigured KMS key. Simplify the parameter names, such as the instance size, by mapping them to Small, Large, and X-Large using the Mappings section in CloudFormation. Display the URL to the notebook using the Outputs section, then upload the template into an AWS Service Catalog product in the data scientist's portfolio, and share it with the data scientist's IAM role.
D. Create an AWS CLI script that the data scientists can run locally. Provide step-by-step instructions about the parameters to be provided while executing the AWS CLI script to launch a Jupyter notebook with a preconfigured KMS key. Distribute the CLI script to the data scientists using a shared Amazon S3 bucket.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 679
A multimedia company needs to deliver its video-on-demand (VOD) content to its subscribers in a cost-effective way. The video files range in size from 1-15 GB and are typically viewed frequently for the first 6 months after creation, and then access decreases considerably. The company requires all video files to remain immediately available for subscribers. There are now roughly 30,000 files, and the company anticipates doubling that number over time.
What is the MOST cost-effective solution for delivering the company's VOD content?

A. Store the video files in an Amazon S3 bucket using S3 Intelligent-Tiering. Use Amazon CloudFront to deliver the content with the S3 bucket as the origin.
B. Use AWS Elemental MediaConvert and store the adaptive bitrate video files in Amazon S3. Configure an AWS Elemental MediaPackage endpoint to deliver the content from Amazon S3.
C. Store the video files in Amazon Elastic File System (Amazon EFS) Standard. Enable EFS lifecycle management to move the video files to EFS Infrequent Access after 6 months. Create an Amazon EC2 Auto Scaling group behind an Elastic Load Balancer to deliver the content from Amazon EFS.
D. Store the video files in Amazon S3 Standard. Create S3 Lifecycle rules to move the video files to S3 Standard-Infrequent Access (S3 Standard-IA) after 6

months and to S3 Glacier Deep Archive after 1 year. Use Amazon CloudFront to deliver the content with the S3 bucket as the origin.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 680
A Development team has created a series of AWS CloudFormation templates to help deploy services. They created a template for a network/virtual private cloud (VPC) stack, a database stack, a bastion host stack, and a web applicationspecific stack. Each service requires the deployment of at least:
A network/VPC stack
A bastion host stack
A web application stack
Each template has multiple input parameters that make it difficult to deploy the services individually from the AWS CloudFormation console. The input parameters from one stack are typically outputs from other stacks. For example, the VPC ID, subnet IDs, and security groups from the network stack may need to be used in the application stack or database stack.
Which actions will help reduce both the operational burden and the number of parameters passed into a service deployment? (Choose two.)

A. Create a new AWS CloudFormation template for each service. Alter the existing templates to use cross-stack references to eliminate passing many parameters to each template. Call each required stack for the application as a nested stack from the new stack. Call the newly created service stack from the AWS CloudFormation console to deploy the specific service with a subset of the parameters previously required.
B. Create a new portfolio in AWS Service Catalog for each service. Create a product for each existing AWS CloudFormation template required to build the service. Add the products to the portfolio that represents that service in AWS Service Catalog.
To deploy the service, select the specific service portfolio and launch the portfolio with the necessary parameters to deploy all templates.
C. Set up an AWS CodePipeline workflow for each service. For each existing template, choose AWS CloudFormation as a deployment action. Add the AWS CloudFormation template to the deployment action. Ensure that the deployment actions are processed to make sure that dependencies are obeyed. Use configuration files and scripts to share parameters between the stacks. To launch the service, execute the specific template by choosing the name of the service and releasing a change.
D. Use AWS Step Functions to define a new service. Create a new AWS CloudFormation template for each service. Alter the existing templates to use cross-stack references to eliminate passing many parameters to each template. Call each required stack for the application as a nested stack from the new service template. Configure AWS Step Functions to call the service template directly. In the AWS Step Functions console, execute the step.
E. Create a new portfolio for the services in AWS Service Catalog. Create a new AWS CloudFormation template for each service. Alter the existing templates to use cross-stack references to eliminate passing many parameters to each template.
Call each required stack for the application as a nested stack from the new stack. Create a product for each application. Add the service template to the product. Add each new product to the portfolio. Deploy the product from the portfolio to deploy the service with the necessary parameters only to start the deployment.

Correct Answer: AE

Section: (none)
Explanation

Explanation/Reference:

QUESTION 681
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 three.)

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 682
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 683
A company has used infrastructure as code (IaC) to provision a set of two Amazon EC2 instances. The instances have remained the same for several years.
The company's business has grown rapidly in the past few months. In response the company's operations team has implemented an Auto Scaling group to manage the sudden increases in traffic. Company policy requires a monthly installation of security updates on all operating systems that are running.
The most recent security update required a reboot. As a result, the Auto Scaling group terminated the instances and replaced them with new, unpatched instances.
Which combination of steps should a solutions architect recommend to avoid a recurrence of this issue? (Choose two.)

A. Modify the Auto Scaling group by setting the Update policy to target the oldest launch configuration for replacement.
B. Create a new Auto Scaling group before the next patch maintenance. During the maintenance window, patch both groups and reboot the instances.
C. Create an Elastic Load Balancer in front of the Auto Scaling group. Configure monitoring to ensure that target group health checks return healthy after the Auto Scaling group replaces the terminated instances.
D. Create automation scripts to patch an AMI, update the launch configuration, and invoke an Auto Scaling instance refresh.
E. Create an Elastic Load Balancer in front of the Auto Scaling group. Configure termination protection on the instances.

Correct Answer: AC
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://medium.com/@endofcake/using-terraform-for-zero-downtime-updates-of-an-auto-scaling-group-in-aws-60faca582664 https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-add-elb-healthcheck.html

QUESTION 684
A solutions architect must analyze a company's Amazon EC2 instances and Amazon Elastic Block Store (Amazon EBS) volumes to determine whether the company is using resources efficiently. The company is running several large, highmemory EC2 instances to host database clusters that are deployed in active/passive configurations. The utilization of these EC2 instances varies by the applications that use the databases, and the company has not identified a pattern.
The solutions architect must analyze the environment and take action based on the findings.
Which solution meets these requirements MOST cost-effectively?

A. Create a dashboard by using AWS Systems Manager OpsCenter. Configure visualizations for Amazon CloudWatch metrics that are associated with the EC2 instances and their EBS volumes. Review the dashboard periodically, and identify usage patterns. Rightsize the EC2 instances based on the peaks in the metrics.
B. Turn on Amazon CloudWatch detailed monitoring for the EC2 instances and their EBS volumes. Create and review a dashboard that is based on the metrics. Identify usage patterns. Rightsize the EC2 instances based on the peaks in the metrics.
C. Install the Amazon CloudWatch agent on each of the EC2 instances. Turn on AWS Compute Optimizer, and let it run for at least 12 hours. Review the recommendations from Compute Optimizer, and rightsize the EC2 instances as directed.
D. Sign up for the AWS Enterprise Support plan. Turn on AWS Trusted Advisor. Wait 12 hours. Review the recommendations from Trusted Advisor, and

rightsize the EC2 instances as directed.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 685
A company is using an on-premises Active Directory service for user authentication. The company wants to use the same authentication service to sign in to the company's AWS accounts, which are using AWS Organizations. AWS Site-to-Site VPN connectivity already exists between the on-premises environment and all the company's AWS accounts.
The company's security policy requires conditional access to the accounts based on user groups and roles. User identities must be managed in a single location.
Which solution will meet these requirements?

A. Configure AWS Single Sign-On (AWS SSO) to connect to Active Directory by using SAML 2.0. Enable automatic provisioning by using the System for Cross-domain Identity Management (SCIM) v2.0 protocol. Grant access to the AWS accounts by using attribute-based access controls (ABACs).
B. Configure AWS Single Sign-On (AWS SSO) by using AWS SSO as an identity source. Enable automatic provisioning by using the System for Cross-domain Identity Management (SCIM) v2.0 protocol. Grant access to the AWS accounts by using AWS SSO permission sets.
C. In one of the company's AWS accounts, configure AWS Identity and Access Management (IAM) to use a SAML 2.0 identity provider. Provision IAM users that are mapped to the federated users. Grant access that corresponds to appropriate groups in Active Directory. Grant access to the required AWS accounts by using cross-account IAM users.
D. In one of the company's AWS accounts, configure AWS Identity and Access Management (IAM) to use an OpenID Connect (OIDC) identity provider. Provision IAM roles that grant access to the AWS account for the federated users that correspond to appropriate groups in Active Directory. Grant access to the required AWS accounts by using cross-account IAM roles.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/singlesignon/latest/userguide/onelogin-idp.html

QUESTION 686
After your Lambda function has been running for some time, you need to look at some metrics to ascertain how your function is performing and decide to use the AWS CLI to do this.
Which of the following commands must be used to access these metrics using the AWS CLI?

A. mon-list-metrics and mon-get-stats
B. list-metrics and get-metric-statistics
C. ListMetrics and GetMetricStatistics
D. list-metrics and mon-get-stats

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
AWS Lambda automatically monitors functions on your behalf, reporting metrics through Amazon CloudWatch.
To access metrics using the AWS CLI
Use the list-metrics and get-metric-statistics commands.
Reference: http://docs.aws.amazon.com/lambda/latest/dg/monitoring-functions-access-metrics.html

QUESTION 687
A company built an ecommerce website on AWS using a three-tier web architecture. The application is Java-based and composed of an Amazon CloudFront distribution, an Apache web server layer of Amazon EC2 instances in an Auto Scaling group, and a backend Amazon Aurora MySQL database.
Last month, during a promotional sales event, users reported errors and timeouts while adding items to their shopping carts.
The operations team recovered the logs created by the web servers and reviewed Aurora DB cluster performance metrics.
Some of the web servers were terminated before logs could be collected and the Aurora metrics were not sufficient for query performance analysis. Which combination of steps must the solutions architect take to improve application performance visibility during peak traffic events? (Choose three.)

A. Configure the Aurora MySQL DB cluster to publish slow query and error logs to Amazon CloudWatch Logs.
B. Implement the AWS X-Ray SDK to trace incoming HTTP requests on the EC2 instances and implement tracing of SQLqueries with the X-Ray SDK for Java.
C. Configure the Aurora MySQL DB cluster to stream slow query and error logs to Amazon Kinesis
D. Install and configure an Amazon CloudWatch Logs agent on the EC2 instances to send the Apache logs to CloudWatch Logs.
E. Enable and configure AWS CloudTrail to collect and analyze application activity from Amazon EC2 and Aurora.
F. Enable Aurora MySQL DB cluster performance benchmarking and publish the stream to AWS X-Ray.

Correct Answer: BCE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 688

When using Numeric Conditions within IAM, short versions of the available comparators can be used instead of the more verbose versions.
Which of the following is the short version of the Numeric Condition "NumericLessThanEquals"?

A. numlteq
B. numlteql
C. numltequals
D. numeql

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
When using Numeric Conditions within IAM, short versions of the available comparators can be used instead of the more verbose versions. For instance, numlteq is the short version of NumericLessThanEquals.
Reference: http://awsdocs.s3.amazonaws.com/SQS/2011-10-01/sqs-dg-2011-10-01.pdf

QUESTION 689
A company wants to move a web application to AWS. The application stores session information locally on each web server, which will make auto scaling difficult. As part of the migration, the application will be rewritten to decouple the session data from the web servers. The company requires low latency, scalability, and availability.
Which service will meet the requirements for storing the session information in the MOST cost-effective way?

A. Amazon ElastiCache with the Memcached engine
B. Amazon S3
C. Amazon RDS MySQL
D. Amazon ElastiCache with the Redis engine

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/caching/session-management/ https://aws.amazon.com/elasticache/redis-vsmemcached/

QUESTION 690
A Solutions Architect must migrate an existing on-premises web application with 70 TB of static files supporting a public open-data initiative. The Architect wants

to upgrade to the latest version of the host operating system as part of the migration effort.
Which is the FASTEST and MOST cost-effective way to perform the migration?

A. Run a physical-to-virtual conversion on the application server. Transfer the server image over the internet, and transfer the static data to Amazon S3.
B. Run a physical-to-virtual conversion on the application server. Transfer the server image over AWS Direct Connect, and transfer the static data to Amazon S3.
C. Re-platform the server to Amazon EC2, and use AWS Snowball to transfer the static data to Amazon S3.
D. Re-platform the server by using the AWS Server Migration Service to move the code and data to a new Amazon EC2 instance.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 691
A company provides auction services for artwork and has users across North America and Europe. The company hosts its application in Amazon EC2 instances in the us-east-1 Region. Artists upload photos of their work as large-size, highresolution image files from their mobile phones to a centralized Amazon S3 bucket created in the us-east-1 Region. The users in Europe are reporting slow performance for their image uploads.
How can a solutions architect improve the performance of the image upload process?

A. Redeploy the application to use S3 multipart uploads.
B. Create an Amazon CloudFront distribution and point to the application as a custom origin.
C. Configure the buckets to use S3 Transfer Acceleration.
D. Create an Auto Scaling group for the EC2 instances and create a scaling policy.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 692
True or false: In a CloudFormation template, you can reuse the same logical ID several times to reference the resources in other parts of the template.

A. True, a logical ID can be used several times to reference the resources in other parts of the template.
B. False, a logical ID must be unique within the template.

C. False, you can mention a resource only once and you cannot reference it in other parts of a template.
D. False, you cannot reference other parts of the template.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In AWS CloudFormation, the logical ID must be alphanumeric (A-Za-z0-9) and unique within the template. You use the logical name to reference the resource in other parts of the template.
Reference: http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/concept-resources.html

QUESTION 693
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 694
A company has many AWS accounts and uses AWS Organizations to manage all of them. A solutions architect must implement a solution that the company can use to share a common network across multiple accounts.
The company's infrastructure team has a dedicated infrastructure account that has a VPC. The infrastructure team must use this account to manage the network. Individual accounts cannot have the ability to manage their own networks. However, individual accounts must be able to create AWS resources within subnets.
Which combination of actions should the solutions architect perform to meet these requirements? (Choose two.)

A. Create a transit gateway in the infrastructure account.
B. Enable resource sharing from the AWS Organizations management account.
C. Create VPCs in each AWS account within the organization in AWS Organizations. Configure the VPCs to share the same CIDR range and subnets as the VPC in the infrastructure account. Peer the VPCs in each individual account with the VPC in the infrastructure account.
D. Create a resource share in AWS Resource Access Manager in the infrastructure account. Select the specific AWS Organizations OU that will use the shared network. Select each subnet to associate with the resource share.
E. Create a resource share in AWS Resource Access Manager in the infrastructure account. Select the specific AWS Organizations OU that will use the shared network. Select each prefix list to associate with the resource share.

Correct Answer: BE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 695
Identify a true statement about the statement ID (Sid) in IAM.

A. You cannot expose the Sid in the IAM API.
B. You cannot use a Sid value as a sub-ID for a policy document's ID for services provided by SQS and SNS.
C. You can expose the Sid in the IAM API.
D. You cannot assign a Sid value to each statement in a statement array.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The Sid (statement ID) is an optional identifier that you provide for the policy statement. You can assign a Sid a value to each statement in a statement array. In IAM, the Sid is not exposed in the IAM API. You can't retrieve a particular statement based on this ID.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements.html#Sid

QUESTION 696
A company's processing team has an AWS account with a production application. The application runs on Amazon EC2 instances behind a Network Load Balancer (NLB). The EC2 instances are hosted in private subnets in a VPC in the eu-west- 1 Region. The VPC was assigned the CIDR block of 10.0.0.0/16. The billing team recently created a new AWS account and deployed an application on EC2 instances that are hosted in private subnets in a VPC in the eu-central-1

Region. The new VPC is assigned the CIDR block of 10.0.0.0/16.
The processing application needs to securely communicate with the billing application over a proprietary TCP port.
What should a solutions architect do to meet this requirement with the LEAST amount of operational effort?

A. In the billing team's account, create a new VPC and subnets in eu-central-1 that use the CIDR block of 192.168.0.0/16. Redeploy the application to the new subnets. Configure a VPC peering connection between the two VPCs.
B. In the processing team's account, add an additional CIDR block of 192.168.0.0/16 to the VPC in eu-west-1. Restart each of the EC2 instances so that they obtain a new IP address. Configure an interRegion VPC peering connection between the two VPCs.
C. In the billing team's account, create a new VPC and subnets in eu-west-1 that use the CIDR block of 192.168.0.0/16. Create a VPC endpoint service (AWS PrivateLink) in the processing team's account and an interface VPC endpoint in the new VPConfigure an inter-Region VPC peering connection in the billing team's account between the two VPCs.
D. In each account, create a new VPC with the CIDR blocks of 192.168.0.0/16 and 172.16.0.0/16. Create inter-Region VPC peering connections between the billing team's VPCs and the processing team's VPCs. Create gateway VPC endpoints to allow traffic to route between the VPCs.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 697
A company is refactoring an existing web service that provides read and write access to structured data. The service must respond to short but significant spikes in the system load. The service must be fault tolerant across multiple AWS Regions.
Which actions should be taken to meet these requirements?

A. Store the data in Amazon DocumentDB. Create a single global Amazon CloudFront distribution with a custom origin built on edge-optimized Amazon API Gateway and AWS Lambda. Assign the company's domain as an alternate domain for the distribution, and configure Amazon Route 53 with an alias to the CloudFront distribution.
B. Store the data in replicated Amazon S3 buckets in two Regions. Create an Amazon CloudFront distribution in each Region, with custom origins built on Amazon API Gateway and AWS Lambda launched in each Region. Assign the company's domain as an alternate domain for both distributions, and configure Amazon Route 53 with a failover routing policy between them.
C. Store the data in an Amazon DynamoDB global table in two Regions using on-demand capacity mode. In both Regions, run the web service as Amazon ECS Fargate tasks in an Auto Scaling ECS service behind an Application Load Balancer (ALB). In Amazon Route 53, configure an alias record in the company's domain and a Route 53 latency-based routing policy with health checks to distribute traffic between the two ALBs.
D. Store the data in Amazon Aurora global databases. Add Auto Scaling replicas to both Regions. Run the web service on Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer in each Region. Configure the instances to download the web service code in the user data. In Amazon Route 53, configure an alias record for the company's domain and a multi-value routing policy

Correct Answer: A

Section: (none)
Explanation

Explanation/Reference:

QUESTION 698
While assigning a tag to an instance, which of the below mentioned options is not a valid tag key/value pair?

A. ey : "aws" Value:"aws"
B. Key: "aws:name" Value: "instanceAnswer: Aws
C. Key: "Name :aws" Value: "instanceAnswer: Aws
D. Key : "nameAnswer: Aws" Value:"aws:instance

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In Amazon Web Services, to help manage EC2 instances as well their usage in a better way, the user can tag the instances.
The tags are metadata assigned by the user which consists of a key and value. The tag key cannot have a prefix as "aws:", although it can have only "aws".
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html

QUESTION 699
A company requires that all internal application connectivity use private IP addresses. To facilitate this policy, a solutions architect has created interface endpoints to connect to AWS public services. Upon testing, the solutions architect notices that the service names are resolving to public IP addresses, and that internal services cannot connect to the interface endpoints.
Which step should the solutions architect take to resolve this issue?

A. Update the subnet route table with a route to the interface endpoint
B. Enable the private DNS option on the VPC attributes
C. Configure the security group on the interface endpoint to allow connectivity to the AWS services
D. Configure an Amazon Route 53 private hosted zone with a conditional forwarder for the internal application

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 700
A solutions architect is migrating an existing workload to AWS Fargate. The task can only run in a private subnet within the VPC where there is no direct connectivity from outside the system to the application. When the Fargate task is launched, the task fails with the following error: CannotPullContainerError: API error (500): Get https://111122223333.dkr.ecr.us-east-1.amazonaws.com/v2/: net/http:request canceled while waiting for connectionHow should the solutions architect correct this error?

A. Ensure the task is set to ENABLED for the auto-assign public IP setting when launching the task.
B. Ensure the task is set to DISABLED for the auto-assign public IP setting when launching the task. Configure a NAT gateway in the public subnet in the VPC to route requests to the internet.
C. Ensure the task is set to DISABLED for the auto-assign public IP setting when launching the task. Configure a NAT gateway in the private subnet in the VPC to route requests to the internet.
D. Ensure the network mode is set to bridge in the Fargate task definition.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 701
To ensure failover capabilities on an elastic network interface (ENI), what should you use for incoming traffic?

A. A Route53 A record
B. A secondary private IP
C. A secondary public IP
D. A secondary ENI

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
To ensure failover capabilities on an elastic network interface (ENI), consider using a secondary private IP for incoming traffic and if a failure occurs, you can move the interface and/or secondary private IP address to a standby instance.

Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html

QUESTION 702
Do you need to use Amazon Cognito to use the Amazon Mobile Analytics service?

A. No. However, it is recommend by AWS to use Amazon Cognito for security best practices.
B. Yes. You need to use it only if you have IAM root access.
C. No. You cannot use it at all, and you need to use AWS IAM accounts.
D. Yes. It is recommended by AWS to use Amazon Cognito to use Amazon Mobile Analytics service.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
You can initialize Amazon Mobile Analytics using AWS IAM accounts. AWS recommend using Amazon Cognito for security best practices.
Reference: http://aws.amazon.com/mobileanalytics/faqs/

QUESTION 703
A company has a requirement that only allows specially hardened AMIs to be launched into public subnets in a VPC, and for the AMIs to be associated with a specific security group. Allowing non-compliant instances to launch into the public subnet could present a significant security risk if they are allowed to operate.
A mapping of approved AMIs to subnets to security groups exists in an Amazon DynamoDB table in the same AWS account.
The company created an AWS Lambda function that, when invoked, will terminate a given Amazon EC2 instance if the combination of AMI, subnet, and security group are not approved in the DynamoDB table.
What should the Solutions Architect do to MOST quickly mitigate the risk of compliance deviations?

A. Create an Amazon CloudWatch Events rule that matches each time an EC2 instance is launched using one of the allowed AMIs, and associate it with the Lambda function as the target.
B. For the Amazon S3 bucket receiving the AWS CloudTrail logs, create an S3 event notification configuration with a filter to match when logs contain the ec2:RunInstances action, and associate it with the Lambda function as the target.
C. Enable AWS CloudTrail and configure it to stream to an Amazon CloudWatch Logs group. Create a metric filter in CloudWatch to match when the ec2:RunInstances action occurs, and trigger the Lambda function when the metric is greater than 0.
D. Create an Amazon CloudWatch Events rule that matches each time an EC2 instance is launched, and associate it with the Lambda function as the target.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 704
A company prefers to limit running Amazon EC2 instances to those that were launched from AMIs pre-approved by the Information Security department. The Development team has an agile continuous integration and deployment process that cannot be stalled by the solution.
Which method enforces the required controls with the LEAST impact on the development process? (Choose two.)

A. Use IAM policies to restrict the ability of users or other automated entities to launch EC2 instances based on a specific set of pre-approved AMIs, such as those tagged in a specific way by Information Security.
B. Use regular scans within Amazon Inspector with a custom assessment template to determine if the EC2 instance that the Amazon Inspector Agent is running on is based upon a pre-approved AMI. If it is not, shut down the instance and inform Information Security by email that this occurred.
C. Only allow launching of EC2 instances using a centralized DevOps team, which is given work packages via notifications from an internal ticketing system. Users make requests for resources using this ticketing tool, which has manual information security approval steps to ensure that EC2 instances are only launched from approved AMIs.
D. Use AWS Config rules to spot any launches of EC2 instances based on non-approved AMIs, trigger an AWS Lambda function to automatically terminate the instance, and publish a message to an Amazon SNS topic to inform Information Security that this occurred.
E. Use a scheduled AWS Lambda function to scan through the list of running instances within the virtual private cloud (VPC) and determine if any of these are based on unapproved AMIs. Publish a message to an SNS topic to inform Information Security that this occurred and then shut down the instance.

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://docs.aws.amazon.com/config/latest/developerguide/evaluate-config_develop-rules_getting-started.html

QUESTION 705
A company uses Amazon S3 to host a web application. Currently, the company uses a continuous integration tool running on an Amazon EC2 instance that builds and deploys the application by uploading it to an S3 bucket. A Solutions Architect needs to enhance the security of the company's platform with the following requirements:
A build process should be run in a separate account from the account hosting the web application. A build process should have minimal access in the account it operates in. Long-lived credentials should not be used.
As a start, the Development team created two AWS accounts: one for the application named web account process; other is a named build account.
Which solution should the Solutions Architect use to meet the security requirements?

A. In the build account, create a new IAM role, which can be assumed by Amazon EC2 only. Attach the role to the EC2 instance running the continuous integration process. Create an IAM policy to allow s3: PutObject calls on the S3 bucket in the web account. In the web account, create an S3 bucket policy

attached to the S3 bucket that allows the build account to use s3:PutObject calls.

B. In the build account, create a new IAM role, which can be assumed by Amazon EC2 only. Attach the role to the EC2 instance running the continuous integration process. Create an IAM policy to allow s3: PutObject calls on the S3 bucket in the web account. In the web account, create an S3 bucket policy attached to the S3 bucket that allows the newly created IAM role to use s3:PutObject calls.

C. In the build account, create a new IAM user. Store the access key and secret access key in AWS Secrets Manager.
Modify the continuous integration process to perform a lookup of the IAM user credentials from Secrets Manager. Create an IAM policy to allow s3: PutObject calls on the S3 bucket in the web account, and attack it to the user. In the web account, create an S3 bucket policy attached to the S3 bucket that allows the newly created IAM user to use s3:PutObject calls.

D. In the build account, modify the continuous integration process to perform a lookup of the IAM user credentials from AWS Secrets Manager. In the web account, create a new IAM user. Store the access key and secret access key in Secrets Manager. Attach the PowerUserAccess IAM policy to the IAM user.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 706
A user is accessing an EC2 instance on the SSH port for IP 10.20.30.40/32.
Which one is a secure way to configure that the instance can be accessed only from this IP?

A. In the security group, open port 22 for IP 10.20.30.40

B. In the security group, open port 22 for IP 10.20.30.0

C. In the security group, open port 22 for IP 10.20.30.40/32

D. In the security group, open port 22 for IP 10.20.30.40/0

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In AWS EC2, while configuring a security group, the user needs to specify the IP address in CIDR notation. The CIDR IP range 10.20.30.40/32 says it is for a single IP 10.20.30.40. If the user specifies the IP as 10.20.30.40 only, the security group will not accept and ask for it in a CIDR format.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html

QUESTION 707
What is the maximum write throughput I can provision for a single Dynamic DB table?

A. 1,000 write capacity units
B. 100,000 write capacity units
C. Dynamic DB is designed to scale without limits, but if you go beyond 10,000 you have to contact AWS first.
D. 10,000 write capacity units

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/dynamodb/faqs/

QUESTION 708
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. The access policy attached to the role limits what user in Account C can do when the user accesses Account A.
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 crossaccount 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:

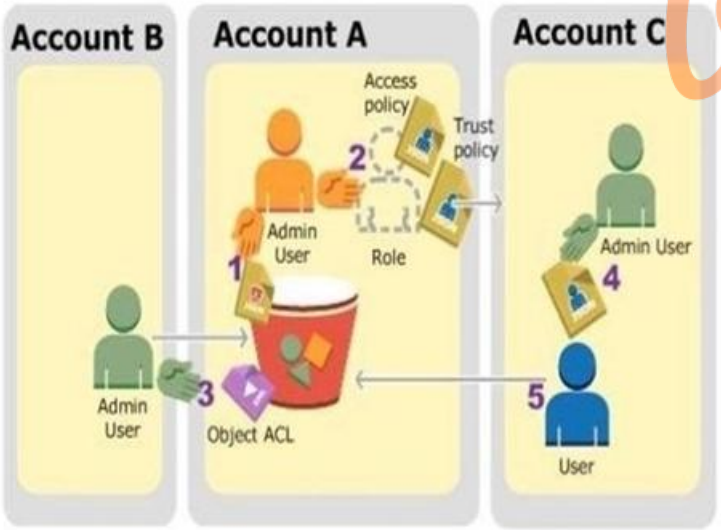

Account A administrator user attaches a bucket policy granting Account B conditional permission to upload objects.
Account A administrator creates an IAM role, establishing trust with Account C, so users in that account can access Account
A. The access policy attached to the role limits what user in 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 709
A company has a data center that must be migrated to AWS as quickly as possible. The data center has a 500 Mbps AWS Direct Connect link and a separate, fully available 1 Gbps ISP connection. A Solutions Architect must transfer 20 TB of data from the data center to an Amazon S3 bucket.
What is the FASTEST way transfer the data?

A. Upload the data to the S3 bucket using the existing DX link.
B. Send the data to AWS using the AWS Import/Export service.
C. Upload the data using an 80 TB AWS Snowball device.
D. Upload the data to the S3 bucket using S3 Transfer Acceleration.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Import/Export supports importing and exporting data into and out of Amazon S3 buckets. For significant data sets, AWS Import/Export is often faster than Internet transfer and more cost effective than upgrading your connectivity.
Reference: https://stackshare.io/stackups/aws-direct-connect-vs-aws-import-export

QUESTION 710
A company wants to migrate an application to Amazon EC2 from VMware Infrastructure that runs in an on-premises data center. A solutions architect must

preserve the software and configuration settings during the migration.
What should the solutions architect do to meet these requirements?

A. Configure the AWS DataSync agent to start replicating the data store to Amazon FSx for Windows File Server. Use the SMB share to host the VMware data store. Use VM Import/Export to move the VMs to Amazon EC2.
B. Use the VMware vSphere client to export the application as an image in Open Virtualization Format (OVF) format. Create an Amazon S3 bucket to store the image in the destination AWS Region. Create and apply an IAM role for VM Import. Use the AWS CLI to run the EC2 import command.
C. Configure AWS Storage Gateway for files service to export a Common Internet File System (CIFS) share. Create a backup copy to the shared folder. Sign in to the AWS Management Console and create an AMI from the backup copy.
Launch an EC2 instance that is based on the AMI.
D. Create a managed-instance activation for a hybrid environment in AWS Systems Manager. Download and install Systems Manager Agent on the on-premises VM. Register the VM with Systems Manager to be a managed instance. Use AWS Backup to create a snapshot of the VM and create an AMI. Launch an EC2 instance that is based on the AMI.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/systems-manager/latest/userguide/sysman-install-managed-linux.html

QUESTION 711
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:
Explanation:
Elastic Beanstalk provides support for running Amazon RDS instances in your Elastic Beanstalk environment. This works great for development and testing environments, but is not ideal for a production environment because it ties the lifecycle of the database instance to the lifecycle of your application's environment. Reference: http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/AWSHowTo.RDS.html

QUESTION 712
Which of the following components of AWS Data Pipeline specifies the business logic of your data management?

A. Task Runner
B. Pipeline definition
C. AWS Direct Connect
D. Amazon Simple Storage Service 9Amazon S3)

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A pipeline definition specifies the business logic of your data management.
Reference: http://docs.aws.amazon.com/datapipeline/latest/DeveloperGuide/what-is-datapipeline.html

QUESTION 713
An organization (account ID 123412341234) has configured the IAM policy to allow the user to modify his credentials.
What will the below mentioned statement allow the user to perform?

```
{
"Version": "2012-10-17",
"Statement": [{
"Effect": "Allow",
"Action": [
"iam:AddUserToGroup",
"iam:RemoveUserFromGroup",
"iam:GetGroup"
],
"Resource": "arn:aws:iam:: 123412341234:group/TestingGroup"
}
]
```

A. Allow the IAM user to update the membership of the group called TestingGroup
B. The IAM policy will throw an error due to an invalid resource name
C. The IAM policy will allow the user to subscribe to any IAM group
D. Allow the IAM user to delete the TestingGroup

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If the organization (account ID 123412341234) wants their users to manage their subscription to the groups, they should create a relevant policy for that. The below mentioned policy allows the respective IAM user to update the membership of the group called MarketingGroup.
{
"Version": "2012-10-17",
"Statement": [{
"Effect": "Allow", "Action": ["iam:AddUserToGroup",
"iam:RemoveUserFromGroup", "iam:GetGroup"
],
"Resource": "arn:aws:iam:: 123412341234:group/ TestingGroup " }]
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/Credentials-Permissions-examples.html#creds-policies-credentials

QUESTION 714
An enterprise runs 103 line-of-business applications on virtual machines in an on-premises data center. Many of the applications are simple PHP, Java, or Ruby web applications, are no longer actively developed, and serve little traffic.

Which approach should be used to migrate these applications to AWS with the LOWEST infrastructure costs?

A. Deploy the applications to single-instance AWS Elastic Beanstalk environments without a load balancer.
B. Use AWS SMS to create AMIs for each virtual machine and run them in Amazon EC2.
C. Convert each application to a Docker image and deploy to a small Amazon ECS cluster behind an Application Load Balancer.
D. Use VM Import/Export to create AMIs for each virtual machine and run them in single-instance AWS Elastic Beanstalk environments by configuring a custom image.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 715
A company is currently in the design phase of an application that will need an RPO of less than 5 minutes and an RTO of less than 10 minutes. The solutions architecture team is forecasting that the database will store approximately 10 TB of data.
As part of the design, they are looking for a database solution that will provide the company with the ability to fail over to a secondary Region.
Which solution will meet these business requirements at the LOWEST cost?

A. Deploy an Amazon Aurora DB cluster and take snapshots of the cluster every 5 minutes. Once a snapshot is complete, copy the snapshot to a secondary Region to serve as a backup in the event of a failure.
B. Deploy an Amazon RDS instance with a cross-Region read replica in a secondary Region. In the event of a failure, promote the read replica to become the primary.
C. Deploy an Amazon Aurora DB cluster in the primary Region and another in a secondary Region. Use AWS DMS to keep the secondary Region in sync.
D. Deploy an Amazon RDS instance with a read replica in the same Region. In the event of a failure, promote the read replica to become the primary.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 716
A company has developed a web application. The company is hosting the application on a group of Amazon EC2 instances behind an Application Load Balancer. The company wants to improve the security posture of the application and plans to use AWS WAF web ACLs. The solution must not adversely affect legitimate traffic to the application.

How should a solutions architect configure the web ACLs to meet these requirements?

A. Set the action of the web ACL rules to Count. Enable AWS WAF logging. Analyze the requests for false positives. Modify the rules to avoid any false positive. Over time, change the action of the web ACL rules from Count to Block.
B. Use only rate-based rules in the web ACLs, and set the throttle limit as high as possible. Temporarily block all requests that exceed the limit. Define nested rules to narrow the scope of the rate tracking.
C. Set the action of the web ACL rules to Block. Use only AWS managed rule groups in the web ACLs. Evaluate the rule groups by using Amazon CloudWatch metrics with AWS WAF sampled requests or AWS WAF logs.
D. Use only custom rule groups in the web ACLs, and set the action to Allow. Enable AWS WAF logging. Analyze the requests for false positives. Modify the rules to avoid any false positive. Over time, change the action of the web ACL rules from Allow to Block.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/waf/latest/developerguide/web-acl-rule-group-settings.html

QUESTION 717
While implementing the policy keys in AWS Direct Connect, if you use and the request comes from an Amazon EC2 instance, the instance's public IP address is evaluated to determine if access is allowed.

A. aws:SecureTransport
B. aws:EpochIP
C. aws:SourceIp
D. aws:CurrentTime

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
While implementing the policy keys in Amazon RDS, if you use aws: SourceIp and the request comes from an Amazon EC2 instance, the instance's public IP address is evaluated to determine if access is allowed.
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/using_iam.html

QUESTION 718

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 blockbased 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: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Volume gateway provides an iSCSI target, which enables you to create volumes and mount them as iSCSI devices from your on-premises application servers.
The volume gateway runs in either a cached or stored mode.
In the cached mode, your primary data is written to S3, while you retain some portion of it locally in a cache for frequently accessed data.
In the stored mode, your primary data is stored locally and your entire dataset is available for low-latency access while asynchronously backed up to AWS.
In either mode, you can take point-in-time snapshots of your volumes and store them in Amazon S3, enabling you to make space-efficient versioned copies of your volumes for data protection and various data reuse needs.

QUESTION 719
A company manages more than 200 separate internet-facing web applications. All of the applications are deployed to AWS in a single AWS Region. The fully qualified domain names (FQDNs) of all of the applications are made available through
HTTPS using Application Load Balancers (ALBs). The ALBs are configured to use public SSL/TLS certificates.
A Solutions Architect needs to migrate the web applications to a multi-region architecture. All HTTPS services shouldcontinue to work without interruption.
Which approach meets these requirements?

A. Request a certificate for each FQDN using AWS KMS. Associate the certificates with the ALBs in the primary AWS Region. Enable cross-region availability in AWS KMS for the certificates and associate the certificates with the ALBs in the secondary AWS Region.
B. Generate the key pairs and certificate requests for each FQDN using AWS KMS. Associate the certificates with the ALBs in both the primary and secondary AWS Regions.
C. Request a certificate for each FQDN using AWS Certificate Manager. Associate the certificates with the ALBs in both the primary and secondary AWS Regions.
D. Request certificates for each FQDN in both the primary and secondary AWS Regions using AWS Certificate Manager.
Associate the certificates with the corresponding ALBs in each AWS Region.

Correct Answer: D

Section: (none)
Explanation

Explanation/Reference:
Explanation:
Certificates in ACM are regional resources. To use a certificate with Elastic Load Balancing for the same fully qualified domain name (FQDN) or set of FQDNs in more than one AWS region, you must request or import a certificate for each region. For certificates provided by ACM, this means you must revalidate each domain name in the certificate for each region. You cannot copy a certificate between regions.
Reference: https://docs.aws.amazon.com/acm/latest/userguide/acm-regions.html

QUESTION 720
A media company uses Amazon DynamoDB to store metadata for its catalog of movies that are available to stream. Each media item contains user-facing content that includes a description of the media, a list of searchable tags, and other similar data. In addition, media items include a list of Amazon S3 key names that relate to movie files. The company stores these movie files in a single S3 bucket that has versioning enabled. The company uses Amazon CloudFront to serve these movie files.
The company has 100,000 media items, and each media item can have many different S3 objects that represent different encodings of the same media. S3 objects that belong to the same media item are grouped together under the same key prefix, which is a random unique ID.
Because of an expiring contract with a media provider, the company must remove 2,000 media items. The company must completely delete all DynamoDB keys and movie files on Amazon S3 that are related to these media items within 36 hours.
The company must ensure that the content cannot be recovered.
Which combination of actions will meet these requirements? (Choose two.)

A. Configure the DynamoDB table with a TTL field. Create and invoke an AWS Lambda function to perform a conditional update. Set the TTL field to the time of the contract's expiration on every affected media item.
B. Configure an S3 Lifecycle object expiration rule that is based on the contract's expiration date.
C. Write a script to perform a conditional delete on all the affected DynamoDB records.
D. Temporarily suspend versioning on the S3 bucket. Create and invoke an AWS Lambda function that deletes affected objects. Reactivate versioning when the operation is complete.
E. Write a script to delete objects from Amazon S3. Specify in each request a NoncurrentVersionExpiration property with a NoncurrentDays attribute set to 0.

Correct Answer: BD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 721
A company that runs applications on AWS recently subscribed to a new software-as-a-service (SaaS) data vendor. The vendor provides the data by way of a REST API that the vendor hosts in its AWS environment. The vendor offers multiple options for connectivity to the API and is working with the company to find

the best way to connect.
The company's AWS account does not allow outbound internet access from its AWS environment. The vendor's services run on AWS in the same Region as the company's applications.
A solutions architect must implement connectivity to the vendor's API so that the API is highly available in the company's VPC.
Which solution will meet these requirements?

A. Connect to the vendor's public API address for the data service
B. Connect to the vendor by way of a VPC peering connection between the vendor's VPC and the company's VPC
C. Connect to the vendor by way of a VPC endpoint service that uses AWS PrivateLink
D. Connect to a public bastion host that the vendor provides. Tunnel the API traffic

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.oracle.com/en-us/iaas/big-data/doc/use-bastion-host-connect-your-service.html

QUESTION 722
A user has created a VPC with CIDR 20.0.0.0/16. The user has created one subnet with CIDR 20.0.0.0/16 by mistake. The user is trying to create another subnet of CIDR 20.0.1.0/24.
How can the user create the second subnet?

A. The user can modify the first subnet CIDR with AWS CLI
B. The user can modify the first subnet CIDR from the console
C. There is no need to update the subnet as VPC automatically adjusts the CIDR of the first subnet based on the second subnet's CIDR
D. It is not possible to create a second subnet with overlapping IP CIDR without deleting the first subnet.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside the subnet. The user can create a subnet with the same size of VPC. However, he cannot create any other subnet since the CIDR of the second subnet will conflict with the first subnet. The user cannot modify the CIDR of a subnet once it is created. Thus, in this case if required, the user has to delete the subnet and create new subnets.

Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

QUESTION 723
Over which of the following Ethernet standards does AWS Direct Connect link your internal network to an AWS Direct Connect location?

A. Single mode fiber-optic cable
B. Multi-mode fiber-optic cable
C. Shielded balanced copper cable
D. Twisted pair cable

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
AWS Direct Connect links your internal network to an AWS Direct Connect location over a standard 1 gigabit or 10 gigabit Ethernet single mode fiber-optic cable.
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.html

QUESTION 724
ABC has created a multi-tenant Learning Management System (LMS). The application is hosted for five different tenants (clients) in the VPCs of the respective AWS accounts of the tenant. ABC wants to setup a centralized server which can connect with the LMS of each tenant upgrade if required. ABC also wants to ensure that one tenant VPC should not be able to connect to the other tenant VPC for security reasons.
How can ABC setup this scenario?

A. ABC has to setup one centralized VPC which will peer in to all the other VPCs of the tenants.
B. ABC should setup VPC peering with all the VPCs peering each other but block the IPs from CIDR of the tenant VPCs to deny them.
C. ABC should setup all the VPCs with the same CIDR but have a centralized VPThis way only the centralized VPC can talk to the other VPCs using VPC peering.
D. ABC should setup all the VPCs meshed together with VPC peering for all VPCs.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that

the user has defined. A VPC peering connection allows the user to route traffic between the peer VPCs using private IP addresses as if they are a part of the same network.
This is helpful when one VPC from the same or different AWS account wants to connect with resources of the other VPC.
The organization wants to setup that one VPC can connect with all the other VPCs but all other VPCs cannot connect among each other. This can be achieved by configuring VPC peering where one VPC is peered with all the other VPCs, but the other VPCs are not peered to each other.
The VPCs are in the same or a separate AWS account and should not have overlapping CIDR blocks.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/PeeringGuide/peering-configurations-full-access.html#many-vpcs-full-acces

QUESTION 725
Within the IAM service a GROUP is regarded as a:

A. A collection of AWS accounts
B. It's the group of EC2 machines that gain the permissions specified in the GROUP.
C. A collection of users.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Use groups to assign permissions to IAM users
Instead of defining permissions for individual IAM users, it's usually more convenient to create groups that relate to job functions (administrators, developers, accounting, etc.), define the relevant permissions for each group, and then assign IAM users to those groups. All the users in an IAM group inherit the permissions assigned to the group. That way, you can make changes for everyone in a group in just one place. As people move around in your company, you can simply change what IAM group their IAM user belongs to.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html#use-groups-for-permissions

QUESTION 726
A company has implemented an ordering system using an event driven architecture. During initial testing, the system stopped processing orders. Further log analysis revealed that one order message in an Amazon Simple Queue Service (Amazon SQS) standard queue was causing an error on the backend and blocking all subsequent order messages. The visibility timeout of the queue is set to 30 seconds, and the backend processing timeout is set to 10 seconds. A solutions architect needs to analyze faulty order messages and ensure that the system continues to process subsequent messages.
Which step should the solutions architect take to meet these requirements?

A. Increase the backend processing timeout to 30 seconds to match the visibility timeout.
B. Reduce the visibility timeout of the queue to automatically remove the faulty message.
C. Configure a new SQS FIFO queue as a dead-letter queue to isolate the faulty messages.
D. Configure a new SQS standard queue as a dead-letter queue to isolate the faulty messages.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/blogs/compute/using-amazon-sqs-dead-letter-queues-to-control-message-failure/

QUESTION 727
A utility company wants to collect usage data every 5 minutes from its smart meters to facilitate time-of-use metering. When a meter sends data to AWS, the data is sent to Amazon API Gateway, processed by an AWS Lambda function and stored in an Amazon DynamoDB table. During the pilot phase, the Lambda functions took from 3 to 5 seconds to complete.
As more smart meters are deployed, the Engineers notice the Lambda functions are taking from 1 to 2 minutes to complete.
The functions are also increasing in duration as new types of metrics are collected from the devices. There are many ProvisionedThroughputExceededException errors while performing PUT operations on DynamoDB, and there are also many TooManyRequestsException errors from Lambda.
Which combination of changes will resolve these issues? (Choose two.)

A. Increase the write capacity units to the DynamoDB table.
B. Increase the memory available to the Lambda functions.
C. Increase the payload size from the smart meters to send more data.
D. Stream the data into an Amazon Kinesis data stream from API Gateway and process the data in batches.
E. Collect data in an Amazon SQS FIFO queue, which triggers a Lambda function to process each message.

Correct Answer: AB
Section: (none)
Explanation

Explanation/Reference:

QUESTION 728
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 three.)

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/UsingKMSEncryption.html

QUESTION 729
The user has provisioned the PIOPS volume with an EBS optimized instance.
Generally speaking, in which I/O chunk should the bandwidth experienced by the user be measured by AWS?

A. 128 KB
B. 256 KB
C. 64 KB
D. 32 KB

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
IOPS are input/output operations per second. Amazon EBS measures each I/O operation per second (that is 256 KB or smaller) as one IOPS.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-io-characteristics.html

QUESTION 730
A user has created a MySQL RDS instance with PIOPS. Which of the below mentioned statements will help user understand the advantage of PIOPS?

A. The user can achieve additional dedicated capacity for the EBS I/O with an enhanced RDS option
B. It uses a standard EBS volume with optimized configuration the stacks
C. It uses optimized EBS volumes and optimized configuration stacks
D. It provides a dedicated network bandwidth between EBS and RDS

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
RDS DB instance storage comes in two types: standard and provisioned IOPS. Standard storage is allocated on the Amazon EBS volumes and connected to the user's DB instance. Provisioned IOPS uses optimized EBS volumes and an optimized configuration stack. It provides additional, dedicated capacity for the EBS I/O.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Welcome.html

QUESTION 731
In the context of policies and permissions in AWS IAM, the Condition element is _____.

A. crucial while writing the IAM policies
B. an optional element
C. always set to null
D. a mandatory element

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The Condition element (or Condition block) lets you specify conditions for when a policy is in effect. The Condition element is optional.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_ElementDescriptions.html

QUESTION 732
A solutions architect is designing a publicly accessible web application that is on an Amazon CloudFront distribution with an Amazon S3 website endpoint as the origin. When the solution is deployed, the website returns an Error 403: Access Denied message.
Which steps should the solutions architect take to correct the issue? (Choose two.)

A. Remove the S3 block public access option from the S3 bucket.
B. Remove the requester pays option from the S3 bucket.
C. Remove the origin access identity (OAI) from the CloudFront distribution.
D. Change the storage class from S3 Standard to S3 One Zone-Infrequent Access (S3 One Zone-IA).
E. Disable S3 object versioning.

Correct Answer: AC
Section: (none)
Explanation

Explanation/Reference:

QUESTION 733
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:
Explanation:
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-metadata.html

QUESTION 734
A company is building an application on AWS. The application sends logs to an Amazon Elasticsearch Service (Amazon ES) cluster for analysis. All data must be stored within a VPC.
Some of the company's developers work from home. Other developers work from three different company office locations.
The developers need to access Amazon ES to analyze and visualize logs directly from their local development machines.
Which solution will meet these requirements?

A. Configure and set up an AWS Client VPN endpoint. Associate the Client VPN endpoint with a subnet in the VPC.
Configure a Client VPN self-service portal. Instruct the developers to connect by using the client for Client VPN.
B. Create a transit gateway, and connect it to the VPC. Create an AWS Site-to-Site VPN. Create an attachment to the transit gateway. Instruct the developers to connect by using an OpenVPN client.
C. Create a transit gateway, and connect it to the VPOrder an AWS Direct Connect connection. Set up a public VIF on the Direct Connect connection. Associate the public VIF with the transit gateway. Instruct the developers to connect to the Direct Connect connection
D. Create and configure a bastion host in a public subnet of the VPC. Configure the bastion host security group to allow SSH access from the company CIDR

ranges. Instruct the developers to connect by using SSH.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/vpn/latest/clientvpn-admin/cvpn-getting-started.html

QUESTION 735
An organization has developed an application which provides a smarter shopping experience. They need to show a demonstration to various stakeholders who may not be able to access the in premise application so they decide to host a demo version of the application on AWS.
Consequently, they will need a fixed elastic IP attached automatically to the instance when it is launched.
In this scenario which of the below mentioned options will not help assign the elastic IP automatically?

A. Write a script which will fetch the instance metadata on system boot and assign the public IP using that metadata.
B. Provide an elastic IP in the user data and setup a bootstrapping script which will fetch that elastic IP and assign it to the instance.
C. Create a controlling application which launches the instance and assigns the elastic IP based on the parameter provided when that instance is booted.
D. Launch instance with VPC and assign an elastic IP to the primary network interface.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
EC2 allows the user to launch On-Demand instances. If the organization is using an application temporarily only for demo purposes the best way to assign an elastic IP would be: Launch an instance with a VPC and assign an EIP to the primary network interface. This way on every instance start it will have the same IP Create a bootstrapping script and provide it some metadata, such as user data which can be used to assign an EIP Create a controller instance which can schedule the start and stop of the instance and provide an EIP as a parameter so that the controller instance can check the instance boot and assign an EIP The instance metadata gives the current instance data, such as the public/private IP. It can be of no use for assigning an EIP.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AESDG-chapter-instancedata.html

QUESTION 736
A financial company needs to create a separate AWS account for a new digital wallet application. The company uses AWS Organizations to manage its accounts. A solutions architect uses the IAM user Support1 from the master account to create a new member account with finance1@example.com as the email address.
What should the solutions architect do to create IAM users in the new member account?

A. Sign in to the AWS Management Console with AWS account root user credentials by using the 64-character password from the initial AWS Organizations email sent to finance1@example.com. Set up the IAM users as required.
B. From the master account, switch roles to assume the OrganizationAccountAccessRole role with the account ID of the new member account. Set up the IAM users as required.
C. Go to the AWS Management Console sign-in page. Choose “Sign in using root account credentials.” Sign in by using the email address finance1@example.com and the master account’s root password. Set up the IAM users as required.
D. Go to the AWS Management Console sign-in page. Sign in by using the account ID of the new member account and the Support1 IAM credentials. Set up the IAM users as required.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_create.html

QUESTION 737
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: D
Section: (none)
Explanation

Explanation/Reference:

Explanation:
Q: Can a global secondary index key be defined on non-unique attributes?
Yes. Unlike the primary key on a table, a GSI index does not require the indexed attributes to be unique. Q: Are GSI key attributes required in all items of a DynamoDB table?
No. GSIs are sparse indexes. Unlike the requirement of having a primary key, an item in a DynamoDB table does not have to contain any of the GSI keys. If a GSI key has both hash and range elements, and a table item omits either of them, then that item will not be indexed by the corresponding GSI. In such cases, a GSI can be very useful in efficiently locating items that have an uncommon attribute.
Reference: https://aws.amazon.com/dynamodb/faqs/

QUESTION 738
What is a circular dependency in AWS CloudFormation?

A. When Nested Stacks depend on each other.
B. When Resources form a Depend On loop.
C. When a Template references an earlier version of itself.
D. When a Template references a region, which references the original Template.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
To resolve a dependency error, add a Depends On attribute to resources that depend on other resources in your template. In some cases, you must explicitly declare dependencies so that AWS CloudFormation can create or delete resources in the correct order. For example, if you create an Elastic IP and a VPC with an Internet gateway in the same stack, the Elastic IP must depend on the Internet gateway attachment. For additional information, see Depends On Attribute.
Reference: http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/troubleshooting.html#troubleshooting-errorsdependency- error

QUESTION 739
A multimedia company with a single AWS account is launching an application for a global user base. The application storage and bandwidth requirements are unpredictable. The application will use Amazon EC2 instances behind an Application Load Balancer as the web tier and will use Amazon DynamoDB as the database tier. The environment for the application must meet the following requirements:
Low latency when accessed from any part of the world
WebSocket support
End-to-end encryption
Protection against the latest security threats
Managed layer 7 DDoS protection
Which actions should the solutions architect take to meet these requirements? (Choose two.)

A. Use Amazon Route 53 and Amazon CloudFront for content distribution. Use Amazon S3 to store static content

B. Use Amazon Route 53 and AWS Transit Gateway for content distribution. Use an Amazon Elastic Block Store (Amazon EBS) volume to store static content
C. Use AWS WAF with AWS Shield Advanced to protect the application
D. Use AWS WAF and Amazon Detective to protect the application
E. Use AWS Shield Standard to protect the application

Correct Answer: BC
Section: (none)
Explanation

Explanation/Reference:

QUESTION 740
A sys admin is maintaining an application on AWS. The application is installed on EC2 and user has configured ELB and Auto Scaling. Considering future load increase, the user is planning to launch new servers proactively so that they get registered with ELB.
How can the user add these instances with Auto Scaling?

A. Decrease the minimum limit of the Auto Scaling group
B. Increase the maximum limit of the Auto Scaling group
C. Launch an instance manually and register it with ELB on the fly
D. Increase the desired capacity of the Auto Scaling group

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A user can increase the desired capacity of the Auto Scaling group and Auto Scaling will launch a new instance as per the new capacity. The newly launched instances will be registered with ELB if Auto Scaling group is configured with ELB. If the user decreases the minimum size the instances will be removed from Auto Scaling. Increasing the maximum size will not add instances but only set the maximum instance cap.
Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/as-manual-scaling.html

QUESTION 741
An organization is planning to host a Wordpress blog as well as joomla CMS on a single instance launched with VPC. The organization wants to create separate domains for each application using Route 53. The organization may have about ten instances each with these two applications. While launching each instance, the organization configured two separate network interfaces (primary + secondary ENI) with their own Elastic IPs to the instance. The suggestion was to use a public IP from AWS instead of an Elastic IP as the number of elastic IPs allocation per region is restricted in the account.
What action will you recommend to the organization?

A. Only Elastic IP can be used by requesting limit increase, since AWS does not assign a public IP to an instance with multiple ENIs.
B. AWS VPC does not attach a public IP to an ENI; so the only way is to use an Elastic IP.
C. I agree with the suggestion but will prefer that the organization should use separate subnets with each ENI for different public IPs.
D. I agree with the suggestion and it is recommended to use a public IP from AWS since the organization is going to use DNS with Route 53.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. An Elastic Network Interface (ENI) is a virtual network interface that the user can attach to an instance in a VPC.
The user can attach up to two ENIs with a single instance. However, AWS cannot assign a public IP when there are two ENIs attached to a single instance. It is recommended to assign an elastic IP in this scenario. If the organization wants more than 5 EIPs they can request AWS to increase the number.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html

QUESTION 742
A company decided to purchase Amazon EC2 Reserved Instances. A solutions architect is tasked with implementing a solution where only the master account in AWS Organizations is able to purchase the Reserved Instances. Current and future member accounts should be blocked from purchasing Reserved Instances.
Which solution will meet these requirements?

A. Create an SCP with the Deny effect on the ec2:PurchaseReservedInstancesOffering action. Attach the SCP to the root of the organization.
B. Create a new organizational unit (OU) Move all current member accounts to the new OU. Create an SCP with the Deny effect on the ec2:PurchaseReservedInstancesOffering action. Attach the SCP to the new OU.
C. Create an AWS Config rule event that triggers automation that will terminate any Reserved Instances launched by member accounts.
D. Create two new organizational units (OUs): OU1 and OU2. Move all member accounts to OU2 and the master account to OU1. Create an SCP with the Allow effect on the ec2:PurchaseReservedInstancesOffering action. Attach the SCP to OU1.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 743
A media company has a static web application that is generated programmatically. The company has a build pipeline that generates HTML content that is

uploaded to an Amazon S3 bucket served by Amazon CloudFront. The build pipeline runs inside a Build Account. The S3 bucket and CloudFront distribution are in a Distribution Account. The build pipeline uploads the files to Amazon S3 using an IAM role in the Build Account. The S3 bucket has a bucket policy that only allows CloudFront to read objects using an origin access identity (OAI). During testing all attempts to access the application using the CloudFront URL result in an HTTP 403 Access Denied response.
What should a solutions architect suggest to the company to allow access the objects in Amazon S3 through CloudFront?

A. Modify the S3 upload process in the Build Account to add the bucket-owner-full-control ACL to the objects at upload.
B. Create a new cross-account IAM role in the Distribution Account with write access to the S3 bucket. Modify the build pipeline to assume this role to upload the files to the Distribution Account.
C. Modify the S3 upload process in the Build Account to set the object owner to the Distribution Account.
D. Create a new IAM role in the Distribution Account with read access to the S3 bucket. Configure CloudFront to use this new role as its OAI. Modify the build pipeline to assume this role when uploading files from the Build Account.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 744
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 745
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 8xlarge 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 8xlarge 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 746
A newspaper organization has an 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 lo 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 lo 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 747
A fleet of Amazon ECS instances is used to poll an Amazon SQS queue and update items in an Amazon DynamoDB database. Items in the table are not being updated, and the SQS queue is filling up. Amazon CloudWatch Logs are showing consistent 400 errors when attempting to update the table. The provisioned write capacity units are appropriately configured, and no throttling is occurring.
What is the LIKELY cause of the failure?

A. The ECS service was deleted.
B. The ECS configuration does not contain an Auto Scaling group.
C. The ECS instance task execution IAM role was modified.
D. The ECS task role was modified.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 748
A mobile app has become very popular, and usage has gone from a few hundred to millions of users. Users capture and upload images of activities within a city, and provide ratings and recommendations. Data access patterns are unpredictable.
The current application is hosted on Amazon EC2 instances behind an Application Load Balancer (ALB). The application is experiencing slowdowns and costs are growing rapidly.
Which changes should a solutions architect make to the application architecture to control costs and improve performance?

A. Create an Amazon CloudFront distribution and place the ALB behind the distribution. Store static content in Amazon S3 in an Infrequent Access storage class.
B. Store static content in an Amazon S3 bucket using the Intelligent Tiering storage class. Use an Amazon CloudFront distribution in front of the S3 bucket and the ALB.
C. Place AWS Global Accelerator in front of the ALB. Migrate the static content to Amazon EFS, and then run an AWS Lambda function to resize the images during the migration process.
D. Move the application code to AWS Fargate containers and swap out the EC2 instances with the Fargate containers.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 749
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 750
A Solutions Architect is designing the storage layer for a data warehousing application. The data files are large, but they have statically placed metadata at the beginning of each file that describes the size and placement of the file's index. The data files are read in by a fleet of Amazon EC2 instances that store the index size, index location, and other category information about the data file in a database. That database is used by Amazon EMR to group files together for deeper analysis.
What would be the MOST cost-effective, high availability storage solution for this workflow?

A. Store the data files in Amazon S3 and use Range GET for each file's metadata, then index the relevant data.
B. Store the data files in Amazon EFS mounted by the EC2 fleet and EMR nodes.
C. Store the data files on Amazon EBS volumes and allow the EC2 fleet and EMR to mount and unmount the volumes where they are needed.
D. Store the content of the data files in Amazon DynamoDB tables with the metadata, index, and data as their own keys.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 751
A user is running a critical batch process which runs for 1 hour and 50 mins every day at a fixed time. Which of the below mentioned options is the right instance type and costing model in this case if the user performs the same task for the whole year?

A. Instance store backed instance with spot instance pricing.
B. EBS backed instance with standard reserved upfront instance pricing.
C. EBS backed scheduled reserved instance with partial instance pricing.
D. EBS backed instance with on-demand instance pricing.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
For Amazon Web Services, the reserved instance (standard or convertible) helps the user save money if the user is going to run the same instance for a longer period. Generally, if the user uses the instances around 30-40% of the year annually it is recommended to use RI. Here as the instance runs only for 1 hour 50 minutes daily, or less than 8 percent of the year, it is not recommended to have RI as it will be costlier. At its highest potential savings, you are still paying 25 percent of an annual cost for a reserved instance you are you using less than 2 hours a day, (or less than 8 percent of each year) you are not saving money.
Even a scheduled reserved instance has a key limitation, which is that it cannot be stopped or rebooted, only manually terminated with a possibility that it could be restarted. You would have to terminate and restart it within the 1 hour 50-minute window, otherwise you would need to wait until the next day. For a critical daily process, this is likely not an option. Spot Instances are not ideal because the process is critical, and must run for a fixed length of time at a fixed time of day. Spot instances would stop and start based on fluctuations in instance pricing, leaving this process potentially unfinished.
The user should use on-demand with EBS in this case. While it has the highest cost, it also has the greatest flexibility to ensure that a critical process like this is always completed.
Reference: http://aws.amazon.com/ec2/purchasing-options/reserved-instances/

QUESTION 752
An organization is setting up an application on AWS to have both High Availability (HA) and Disaster Recovery (DR). The organization wants to have both Recovery point objective (RPO) and Recovery time objective (RTO) of 10 minutes.
Which of the below mentioned service configurations does not help the organization achieve the said RPO and RTO?

A. Take a snapshot of the data every 10 minutes and copy it to the other region.
B. Use an elastic IP to assign to a running instance and use Route 53 to map the user's domain with that IP.
C. Create ELB with multi-region routing to allow automated failover when required.
D. Use an AMI copy to keep the AMI available in other regions.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

Explanation:
AWS provides an on demand, scalable infrastructure. AWS EC2 allows the user to launch On- Demand instances and the organization should create an AMI of the running instance. Copy the AMI to another region to enable Disaster Recovery (DR) in case of region failure. The organization should also use EBS for persistent storage and take a snapshot every 10 minutes to meet Recovery time objective (RTO). They should also setup an elastic IP and use it with Route 53 to route requests to the same IP. When one of the instances fails the organization can launch new instances and assign the same EIP to a new instance to achieve High Availability (HA). The ELB works only for a particular region and does not route requests across regions. Reference: http:// d36cz9buwru1tt.cloudfront.net/AWS_Disaster_Recovery.pdf

QUESTION 753
A financial services company in North America plans to release a new online web application to its customers on AWS. The company will launch the application in the us-east-1 Region on Amazon EC2 instances. The application must be highly available and must dynamically scale to meet user traffic. The company also wants to implement a disaster recovery environment for the application in the us-west1 Region by using active-passive failover.
Which solution will meet these requirements?

A. Create a VPC in us-east-1 and a VPC in us-west-1. Configure VPC peering. In the us-east-1 VPC, create an Application Load Balancer (ALB) that extends across multiple Availability Zones in both VPCs. Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in both VPCs. Place the Auto Scaling group behind the ALB.
B. Create a VPC in us-east-1 and a VPC in us-west-1. In the us-east-1 VPC, create an Application Load Balancer (ALB) that extends across multiple Availability Zones in that VPC. Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in the us-east-1 VPC. Place the Auto Scaling group behind the ALSet up the same configuration in the us-west-1 VPC. Create an Amazon Route 53 hosted zone. Create separate records for each ALEnable health checks to ensure high availability between Regions.
C. Create a VPC in us-east-1 and a VPC in us-west-1. In the us-east-1 VPC, create an Application Load Balancer (ALB) that extends across multiple Availability Zones in that VPCreate an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in the us-east-1 VPPlace the Auto Scaling group behind the ALB. Set up the same configuration in the us-west-1 VPCreate an Amazon Route 53 hosted zone. Create separate records for each ALB. Enable health checks and configure a failover routing policy for each record.
D. Create a VPC in us-east-1 and a VPC in us-west-1. Configure VPC peering. In the us-east-1 VPC, create an Application Load Balancer (ALB) that extends across multiple Availability Zones in both VPCs. Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in both VPCs. Place the Auto Scaling group behind the ALB. Create an Amazon Route 53 hosted zone. Create a record for the ALB.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 754
By default, Amazon Cognito maintains the last-written version of the data. You can override this behavior and resolve data conflicts programmatically.
In addition, push synchronization allows you to use Amazon Cognito to send a silent notification to all devices associated with an identity to notify them that new data is available.

A. get
B. post
C. pull
D. push

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: http://aws.amazon.com/cognito/faqs/

QUESTION 755
A company has released a new version of a website to target an audience in Asia and South America. The website's media assets are hosted on Amazon S3 and have an Amazon CloudFront distribution to improve end-user performance. However, users are having a poor login experience, the authentication service is only available in the us-east-1 AWS Region.
How can the Solutions Architect improve the login experience and maintain high security and performance with minimal management overhead?

A. Replicate the setup in each new geography and use Amazon Route 53 geo-based routing to route traffic to the AWS Region closest to the users.
B. Use an Amazon Route 53 weighted routing policy to route traffic to the CloudFront distribution. Use CloudFront cached HTTP methods to improve the user login experience.
C. Use Amazon Lambda@Edge attached to the CloudFront viewer request trigger to authenticate and authorize users by maintaining a secure cookie token with a session expiry to improve the user experience in multiple geographies.
D. Replicate the setup in each geography and use Network Load Balancers to route traffic to the authentication service running in the closest region to users.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://aws.amazon.com/blogs/networking-and-content-delivery/authorizationedge-how-to-use-lambdaedge-and-json-webtokens-to-enhance-web-application-security/

QUESTION 756
A company has developed a new release of a popular video game and wants to make it available for public download. The new release package is

approximately 5 GB in size. The company provides downloads for existing releases from a Linuxbased, publicly facing FTP site hosted in an on-premises data center. The company expects the new release will be downloaded by users worldwide. The company wants a solution that provides improved download performance and low transfer costs, regardless of a user's location.
Which solutions will meet these requirements?

A. Store the game files on Amazon EBS volumes mounted on Amazon EC2 instances within an Auto Scaling group. Configure an FTP service on the EC2 instances. Use an Application Load Balancer in front of the Auto Scaling group. Publish the game download URL for users to download the package.
B. Store the game files on Amazon EFS volumes that are attached to Amazon EC2 instances within an Auto Scaling group. Configure an FTP service on each of the EC2 instances. Use an Application Load Balancer in front of the Auto Scaling group. Publish the game download URL for users to download the package.
C. Configure Amazon Route 53 and an Amazon S3 bucket for website hosting. Upload the game files to the S3 bucket. Use Amazon CloudFront for the website. Publish the game download URL for users to download the package.
D. Configure Amazon Route 53 and an Amazon S3 bucket for website hosting. Upload the game files to the S3 bucket. Set Requester Pays for the S3 bucket. Publish the game download URL for users to download the package.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 757
A user is trying to create a PIOPS EBS volume with 4000 IOPS and 100 GB size. AWS does not allow the user to create this volume.
What is the possible root cause for this?

A. PIOPS is supported for EBS higher than 500 GB size
B. The maximum IOPS supported by EBS is 3000
C. The ratio between IOPS and the EBS volume is higher than 30
D. The ratio between IOPS and the EBS volume is lower than 50

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 758
An enterprise company is using a multi-account AWS strategy. There are separate accounts for development staging and production workloads. To control costs and improve governance the following requirements have been defined:
The company must be able to calculate the AWS costs for each project.
The company must be able to calculate the AWS costs for each environment development staging and production.
Commonly deployed IT services must be centrally managed.
Business units can deploy pre-approved IT services only.
Usage of AWS resources in the development account must be limited.
Which combination of actions should be taken to meet these requirements? (Choose three.)

A. Apply environment, cost center, and application name tags to all taggable resources.
B. Configure custom budgets and define thresholds using Cost Explorer.
C. Configure AWS Trusted Advisor to obtain weekly emails with cost-saving estimates.
D. Create a portfolio for each business unit and add products to the portfolios using AWS CloudFormation in AWS Service Catalog.
E. Configure a billing alarm in Amazon CloudWatch.
F. Configure SCPs in AWS Organizations to allow services available using AWS.

Correct Answer: CEF
Section: (none)
Explanation

Explanation/Reference:

QUESTION 759
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 4x1TB 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-1a. 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-1a. Create a RAID 0 volume using the four 800GB SSD ephemeral disks provided with the instance. Configure synchronous, block-level replication to an identically configured instance in us-east-1b.

Correct Answer: E
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://acloud.guru/course/aws-certified-solutions-architect-associate/discuss/-KJdi4tFMp2x_O88J6U4/anarchitecture-design-question

QUESTION 760
A company uses AWS Organizations. The company has an organization that has a central management account. The company plans to provision multiple AWS accounts for different departments. All department accounts must be a member of the company's organization.
Compliance requirements state that each account must have only one VPC. Additionally, each VPC must have an identical network security configuration that includes fully configured subnets, gateways, network ACLs, and security groups.
The company wants this security setup to be automatically applied when a new department account is created. The company wants to use the central management account for all security operations, but the central management account should not have the security setup.
Which approach meets these requirements with the LEAST amount of setup?

A. Create an OU within the company's organization. Add department accounts to the OU. From the central management account, create an AWS CloudFormation template that includes the VPC and the network security configurations. Create a CloudFormation stack set by using this template file with automated deployment enabled. Apply the CloudFormation stack set to the OU.
B. Create a new organization with the central management account. Invite all AWS department accounts into the new organization. From the central management account, create an AWS CloudFormation template that includes the VPC and the network security configurations. Create a CloudFormation stack that is based on this template. Apply the CloudFormation stack to the newly created organization.
C. Invite department accounts to the company's organization. From the central management account, create an AWS CloudFormation template that includes the VPC and the network security configurations. Create an AWS CodePipeline pipeline that will deploy the network security setup to the newly created account. Specify the creation of an account as an event hook. Apply the event hook to the pipeline.
D. Invite department accounts to the company's organization. From the central management account, create an AWS CloudFormation template that includes the VPC and the network security configurations. Create an AWS Lambda function that will deploy the VPC and the network security setup to the newly created account. Create an event that watches for account creation. Configure the event to invoke the pipeline.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/blogs/security/how-to-use-aws-organizations-to-automate-end-to-end-account-creation/

QUESTION 761
A healthcare company runs a production workload on AWS that stores highly sensitive personal information. The security team mandates that, for auditing purposes, any AWS API action using AWS account root user credentials must automatically create a high-priority ticket in the company's ticketing system. The ticketing system has a monthly 3-hour maintenance window when no tickets can be created.
To meet security requirements, the company enabled AWS CloudTrail logs and wrote a scheduled AWS Lambda function that uses Amazon Athena to query API actions performed by the root user. The Lambda function submits any actions found to the ticketing system API. During a recent security audit, the security team discovered that several tickets were not created because the ticketing system was unavailable due to planned maintenance.
Which combination of steps should a solutions architect take to ensure that the incidents are reported to the ticketing system even during planned maintenance? (Choose two.)

A. Create an Amazon SNS topic to which Amazon CloudWatch alarms will be published. Configure a CloudWatch alarm to invoke the Lambda function.
B. Create an Amazon SQS queue to which Amazon CloudWatch alarms will be published. Configure a CloudWatch alarm to publish to the SQS queue.
C. Modify the Lambda function to be triggered by messages published to an Amazon SNS topic. Update the existing application code to retry every 5 minutes if the ticketing system's API endpoint is unavailable.
D. Modify the Lambda function to be triggered when there are messages in the Amazon SQS queue and to return successfully when the ticketing system API has processed the request.
E. Create an Amazon EventBridge rule that triggers on all API events where the invoking user identity is root. Configure the EventBridge rule to write the event to an Amazon SQS queue.

Correct Answer: BD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 762
A large company has many business units. Each business unit has multiple AWS accounts for different purposes. The CIO of the company sees that each business unit has data that would be useful to share with other parts of the company. In total, there are about 10 PB of data that needs to be shared with users in 1,000 AWS accounts. The data is proprietary, so some of it should only be available to users with specific job types. Some of the data is used for throughput of intensive workloads, such as simulations. The number of accounts changes frequently because of new initiatives, acquisitions, and divestitures.
A Solutions Architect has been asked to design a system that will allow for sharing data for use in AWS with all of the employees in the company.
Which approach will allow for secure data sharing in scalable way?

A. Store the data in a single Amazon S3 bucket. Create an IAM role for every combination of job type and business unit that allows for appropriate read/write access based on object prefixes in the S3 bucket. The roles should have trust policies that allow the business unit's AWS accounts to assume their roles. Use IAM in each business unit's AWS account to prevent them from assuming roles for a different job type. Users get credentials to access the data by using AssumeRole from their business unit's AWS account. Users can then use those credentials with an S3 client.
B. Store the data in a single Amazon S3 bucket. Write a bucket policy that uses conditions to grant read and write access where appropriate, based on each

user's business unit and job type. Determine the business unit with the AWS account accessing the bucket and the job type with a prefix in the IAM user's name. Users can access data by using IAM credentials from their business unit's AWS account with an S3 client.

C. Store the data in a series of Amazon S3 buckets. Create an application running in Amazon EC2 that is integrated with the company's identity provider (IdP) that authenticates users and allows them to download or upload data through the application. The application uses the business unit and job type information in the IdP to control what users can upload and download through the application. The users can access the data through the application's API.

D. Store the data in a series of Amazon S3 buckets. Create an AWS STS token vending machine that is integrated with the company's identity provider (IdP). When a user logs in, have the token vending machine attach an IAM policy that assumes the role that limits the user's access and/or upload only the data the user is authorized to access. Users can get credentials by authenticating to the token vending machine's website or API and then use those credentials with an S3 client.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 763
A large company has a business-critical application that runs in a single AWS Region. The application consists of multiple Amazon EC2 instances and an Amazon RDS Multi-AZ DB instance. The EC2 instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones.
A solutions architect is implementing a disaster recovery (DR) plan for the application. The solutions architect has created a pilot light application deployment in a new Region, which is referred to as the DR Region. The DR environment has an Auto Scaling group with a single EC2 instance and a read replica of the RDS DB instance.
The solutions architect must automate a failover from the primary application environment to the pilot light environment in the DR Region.
Which solution meets these requirements with the MOST operational efficiency?

A. Publish an application availability metric to Amazon CloudWatch in the DR Region from the application environment in the primary Region. Create a CloudWatch alarm in the DR Region that is invoked when the application availability metric stops being delivered. Configure the CloudWatch alarm to send a notification to an Amazon Simple Notification Service (Amazon SNS) topic in the DR Region. Add an email subscription to the SNS topic that sends messages to the application owner.
Upon notification, instruct a systems operator to sign in to the AWS Management Console and initiate failover operations for the application.

B. Create a cron task that runs every 5 minutes by using one of the application's EC2 instances in the primary Region.
Configure the cron task to check whether the application is available. Upon failure, the cron task notifies a systems operator and attempts to restart the application services.

C. Create a cron task that runs every 5 minutes by using one of the application's EC2 instances in the primary Region.
Configure the cron task to check whether the application is available. Upon failure, the cron task modifies the DR environment by promoting the read replica and by adding EC2 instances to the Auto Scaling group.

D. Publish an application availability metric to Amazon CloudWatch in the DR Region from the application environment in the primary Region. Create a CloudWatch alarm in the DR Region that is invoked when the application availability metric stops being delivered. Configure the CloudWatch alarm to send a notification to an Amazon Simple Notification Service (Amazon SNS) topic in the DR Region. Use an AWS Lambda function that is invoked by Amazon SNS

in the DR Region to promote the read replica and to add EC2 instances to the Auto Scaling group.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 764
A large company experienced a drastic increase in its monthly AWS spend. This is after Developers accidentally launched Amazon EC2 instances in unexpected regions. The company has established practices around least privileges for Developers and controls access to on-premises resources using Active Directory groups. The company now want to control costs by restricting the level of access that Developers have to the AWS Management Console without impacting their productivity. The company would also like to allow Developers to launch Amazon EC2 in only one region, without limiting access to other services in any region.
How can this company achieve these new security requirements while minimizing the administrative burden on the Operations team?

A. Set up SAML-based authentication tied to an IAM role that has an AdministrativeAccess managed policy attached to it. Attach a customer managed policy that denies access to Amazon EC2 in each region except for the one required.
B. Create an IAM user for each Developer and add them to the developer IAM group that has the PowerUserAccess managed policy attached to it. Attach a customer managed policy that allows the Developers access to Amazon EC2 only in the required region.
C. Set up SAML-based authentication tied to an IAM role that has a PowerUserAccess managed policy and a customer managed policy that deny all the Developers access to any AWS services except AWS Service Catalog. Within AWS Service Catalog, create a product containing only the EC2 resources in the approved region.
D. Set up SAML-based authentication tied to an IAM role that has the PowerUserAccess managed policy attached to it. Attach a customer managed policy that denies access to Amazon EC2 in each region except for the one required.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 765
A company has several applications running in an on-premises data center. The data center runs a mix of Windows and Linux VMs managed by VMware vCenter. A solutions architect needs to create a plan to migrate the applications to AWS.
However, the solutions architect discovers that the document for the applications is not up to date and that there are no complete infrastructure diagrams. The company's developers lack time to discuss their applications and current usage with the solutions architect.
What should the solutions architect do to gather the required information?

A. Deploy the AWS Server Migration Service (AWS SMS) connector using the OVA image on the VMware cluster to collect configuration and utilization data from the VMs.
B. Use the AWS Migration Portfolio Assessment (MPA) tool to connect to each of the VMs to collect the configuration and utilization data.
C. Install the AWS Application Discovery Service on each of the VMs to collect the configuration and utilization data.
D. Register the on-premises VMs with the AWS Migration Hub to collect configuration and utilization data.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 766
A company is hosting an image-processing service on AWS in a VPC. The VPC extends across two Availability Zones. Each Availability Zone contains one public subnet and one private subnet.
The service runs on Amazon EC2 instances in the private subnets. An Application Load Balancer in the public subnets is in front of the service. The service needs to communicate with the internet and does so through two NAT gateways. The service uses Amazon S3 for image storage. The EC2 instances retrieve approximately 1 ?? of data from an S3 bucket each day.
The company has promoted the service as highly secure. A solutions architect must reduce cloud expenditures as much as possible without compromising the service's security posture or increasing the time spent on ongoing operations.
Which solution will meet these requirements?

A. Replace the NAT gateways with NAT instances. In the VPC route table, create a route from the private subnets to the NAT instances.
B. Move the EC2 instances to the public subnets. Remove the NAT gateways.
C. Set up an S3 gateway VPC endpoint in the VPAttach an endpoint policy to the endpoint to allow the required actions on the S3 bucket.
D. Attach an Amazon Elastic File System (Amazon EFS) volume to the EC2 instances. Host the image on the EFS volume.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Create Amazon S3 gateway endpoint in the VPC and add a VPC endpoint policy. This VPC endpoint policy will have a statement that allows S3 access only via access points owned by the organization.
Reference: https://lifesciences-resources.awscloud.com/aws-storage-blog/managing-amazon-s3-access-with-vpc-endpointsand-s3-access-points?Languages=Korean

QUESTION 767
How can a user list the IAM Role configured as a part of the launch config?

A. as-describe-launch-configs -iam-profile
B. as-describe-launch-configs -show-long
C. as-describe-launch-configs -iam-role
D. as-describe-launch-configs -role

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
As-describe-launch-configs describes all the launch config parameters created by the AWS account in the specified region.
Generally, it returns values, such as Launch Config name, Instance Type and AMI ID. If the user wants additional parameters, such as the IAM Profile used in the config, he has to run command: as-describe-launch-configs --show-long

QUESTION 768
A company that develops consumer electronics with offices in Europe and Asia has 60 TB of software images stored on premises in Europe. The company wants to transfer the images to an Amazon S3 bucket in the ap-northeast-1 Region. New software images are created daily and must be encrypted in transit. The company needs a solution that does not require custom development to automatically transfer all existing and new software images to Amazon S3.
What is the next step in the transfer process?

A. Deploy an AWS DataSync agent and configure a task to transfer the images to the S3 bucket
B. Configure Amazon Kinesis Data Firehose to transfer the images using S3 Transfer Acceleration
C. Use an AWS Snowball device to transfer the images with the S3 bucket as the target
D. Transfer the images over a Site-to-Site VPN connection using the S3 API with multipart upload

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 769
You set up your first Lambda function and want to set up some Cloudwatch metrics to monitor your function. Which of the following Lambda metrics can Cloudwatch monitor?

A. Total requests only
B. Status Check Failed, total requests, and error rates
C. Total requests and CPU utilization
D. Total invocations, errors, duration, and throttles

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
AWS Lambda automatically monitors functions on your behalf, reporting metrics through Amazon CloudWatch (CloudWatch). These metrics include total invocations, errors, duration, and throttles.
Reference: http://docs.aws.amazon.com/lambda/latest/dg/monitoring-functions-metrics.html

QUESTION 770
A company has multiple business units. Each business unit has its own AWS account and runs a single website within that account. The company also has a single logging account. Logs from each business unit website are aggregated into a single Amazon S3 bucket in the logging account. The S3 bucket policy provides each business unit with access to write data into the bucket and requires data to be encrypted.
The company needs to encrypt logs uploaded into the bucket using a single AWS Key Management Service (AWS KMS) CMK. The CMK that protects the data must be rotated once every 365 days.
Which strategy is the MOST operationally efficient for the company to use to meet these requirements?

A. Create a customer managed CMK in the logging account. Update the CMK key policy to provide access to the logging account only. Manually rotate the CMK every 365 days.
B. Create a customer managed CMK in the logging account. Update the CMK key policy to provide access to the logging account and business unit accounts. Enable automatic rotation of the CMK.
C. Use an AWS managed CMK in the logging account. Update the CMK key policy to provide access to the logging account and business unit accounts. Manually rotate the CMK every 365 days.
D. Use an AWS managed CMK in the logging account. Update the CMK key policy to provide access to the logging account only. Enable automatic rotation of the CMK.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 771
A company has a web application that securely uploads pictures and videos to an Amazon S3 bucket. The company requires that only authenticated users are allowed to post content. The application generates a presigned URL that is used to upload objects through a browser interface. Most users are reporting slow upload times for objects larger than 100 MB.
What can a Solutions Architect do to improve the performance of these uploads while ensuring only authenticated users are allowed to post content?

A. Set up an Amazon API Gateway with an edge-optimized API endpoint that has a resource as an S3 service proxy.
Configure the PUT method for this resource to expose the S3 PutObject operation. Secure the API Gateway using a COGNITO_USER_POOLS authorizer.
Have the browser interface use API Gateway instead of the presigned URL to upload objects.
B. Set up an Amazon API Gateway with a regional API endpoint that has a resource as an S3 service proxy. Configure the PUT method for this resource to expose the S3 PutObject operation. Secure the API Gateway using an AWS Lambda authorizer. Have the browser interface use API Gateway instead of the presigned URL to upload API objects.
C. Enable an S3 Transfer Acceleration endpoint on the S3 bucket. Use the endpoint when generating the presigned URL.
Have the browser interface upload the objects to this URL using the S3 multipart upload API.
D. Configure an Amazon CloudFront distribution for the destination S3 bucket. Enable PUT and POST methods for the CloudFront cache behavior. Update the CloudFront origin to use an origin access identity (OAI). Give the OAI user s3:PutObject permissions in the bucket policy. Have the browser interface upload objects using the CloudFront distribution.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 772
A large company has increased its utilization of AWS over time in an unmanaged way. As such, they have a large number of independent AWS accounts across different business units, projects, and environments. The company has created a Cloud Center of Excellence team, which is responsible for managing all aspects of the AWS Cloud, including their AWS accounts.
Which of the following should the Cloud Center of Excellence team do to BEST address their requirements in a centralized way? (Choose two.)

A. Control all AWS account root user credentials. Assign AWS IAM users in the account of each user who needs to access AWS resources. Follow the policy of least privilege in assigning permissions toeach user.
B. Tag all AWS resources with details about the business unit, project, and environment. Send all AWS Cost and Usage reports to a central Amazon S3 bucket, and use tools such as Amazon Athena and Amazon QuickSight to collect billing details by business unit.
C. Use the AWS Marketplace to choose and deploy a Cost Management tool. Tag all AWS resources with details about the business unit, project, and environment. Send all AWS Cost and Usage reports for the AWS accounts to this tool for analysis.
D. Set up AWS Organizations. Enable consolidated billing, and link all existing AWS accounts to a master billing account.
Tag all AWS resources with details about the business unit, project and environment. Analyze Cost and Usage reports using tools such as Amazon Athena

and Amazon QuickSight, to collect billing details by business unit.

E. Using a master AWS account, create IAM users within the master account. Define IAM roles in the other AWS accounts, which cover each of the required functions in the account. Follow the policy of least privilege in assigning permissions to each role, then enable the IAM users to assume the roles that they need to use.

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 773
One of the AWS account owners faced a major challenge in June as his account was hacked and the hacker deleted all the data from his AWS account. This resulted in a major blow to the business.
Which of the below mentioned steps would not have helped in preventing this action?

A. Setup an MFA for each user as well as for the root account user.
B. Take a backup of the critical data to offsite / on premise.
C. Create an AMI and a snapshot of the data at regular intervals as well as keep a copy to separate regions.
D. Do not share the AWS access and secret access keys with others as well do not store it inside programs, instead use IAM roles.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
AWS security follows the shared security model where the user is as much responsible as Amazon. If the user wants to have secure access to AWS while hosting applications on EC2, the first security rule to follow is to enable MFA for all users. This will add an added security layer. In the second step, the user should never give his access or secret access keys to anyone as well as store inside programs. The better solution is to use IAM roles. For critical data of the organization, the user should keep an offsite/ in premise backup which will help to recover critical data in case of security breach. It is recommended to have AWS AMIs and snapshots as well as keep them at other regions so that they will help in the DR scenario. However, in case of a data security breach of the account they may not be very helpful as hacker can delete that.
Therefore, creating an AMI and a snapshot of the data at regular intervals as well as keep a copy to separate regions, would not have helped in preventing this action.

QUESTION 774
A company is running a two-tier web-based application in an on-premises data center. The application user consists of a single server running a stateful application. The application connects to a PostgreSQL database running on a separate server. The application's user base is expected to grow significantly, so

the company is migrating the application and database to AWS. The solution will use Amazon Aurora PostgreSQL, Amazon EC2 Auto Scaling, and Elastic Load Balancing.
Which solution will provide a consistent user experience that will allow the application and database tiers to scale?

A. Enable Aurora Auto Scaling for Aurora Replicas. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.
B. Enable Aurora Auto Scaling for Aurora writes. Use an Application Load Balancer with the round robin routing algorithm and sticky sessions enabled.
C. Enable Aurora Auto Scaling for Aurora Replicas. Use an Application Load Balancer with the robin routing and sticky sessions enabled.
D. Enable Aurora Scaling for Aurora writers. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 775
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 three.)

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.

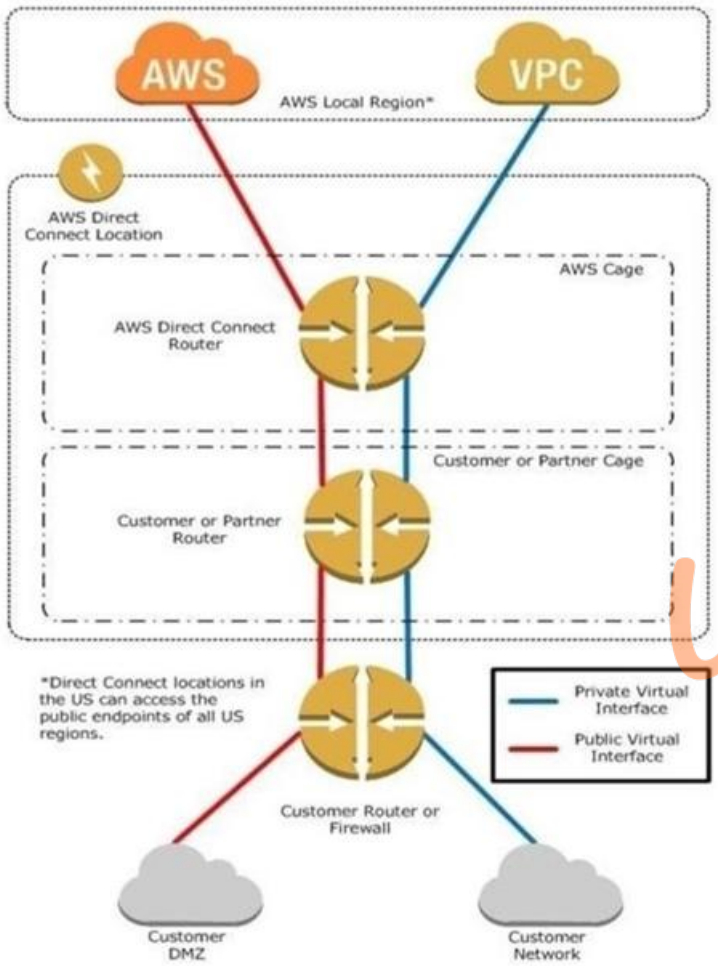

Requirements

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 10GBASELR (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 776
An online retailer needs to regularly process large product catalogs, which are handled in batches. These are sent out to be processed by people using the Amazon Mechanical Turk service, but the retailer has asked its Solutions Architect to design a workflow orchestration system that allows it to handle multiple concurrent Mechanical Turk operations, deal with the result assessment process, and reprocess failures.
Which of the following options gives the retailer the ability to interrogate the state of every workflow with the LEAST amount of implementation effort?

A. Trigger Amazon CloudWatch alarms based upon message visibility in multiple Amazon SQS queues (one queue per workflow stage) and send messages via Amazon SNS to trigger AWS Lambda functions to process the next step. Use Amazon ES and Kibana to visualize Lambda processing logs to see the workflow states.
B. Hold workflow information in an Amazon RDS instance with AWS Lambda functions polling RDS for status changes.
Worker Lambda functions then process the next workflow steps. Amazon QuickSight will visualize workflow states directly out of Amazon RDS.
C. Build the workflow in AWS Step Functions, using it to orchestrate multiple concurrent workflows. The status of each workflow can be visualized in the AWS Management Console, and historical data can be written to Amazon S3 and visualized using Amazon QuickSight.
D. Use Amazon SWF to create a workflow that handles a single batch of catalog records with multiple worker tasks to extract the data, transform it, and send it through Mechanical Turk. Use Amazon ES and Kibana to visualize AWS Lambda processing logs to see the workflow states.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 777
Which of the following is true of an instance profile when an IAM role is created using the console?

A. The instance profile uses a different name.
B. The console gives the instance profile the same name as the role it corresponds to.

C. The instance profile should be created manually by a user.
D. The console creates the role and instance profile as separate actions.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Amazon EC2 uses an instance profile as a container for an IAM role. When you create an IAM role using the console, the console creates an instance profile automatically and gives it the same name as the role it corresponds to. If you use the AWS CLI, API, or an AWS SDK to create a role, you create the role and instance profile as separate actions, and you might give them different names.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_switch-role-ec2_instance-profiles.html

QUESTION 778
A company has a single AWS master billing account, which is the root of the AWS Organizations hierarchy.
The company has multiple AWS accounts within this hierarchy, all organized into organization units (OUs). More OUs and AWS accounts will continue to be created as other parts of the business migrate applications to AWS. These business units may need to use different AWS services. The Security team is implementing the following requirements for all current and future AWS accounts:
Control policies must be applied across all accounts to prohibit AWS servers. Exceptions to the control policies are allowed based on valid use cases.
Which solution will meet these requirements with minimal optional overhead?

A. Use an SCP in Organizations to implement a deny list of AWS servers. Apply this SCP at the level. For any specific exceptions for an OU, create a new SCP for that OU and add the required AWS services to the allow list.
B. Use an SCP in Organizations to implement a deny list of AWS service. Apply this SCP at the root level and each OU.
Remove the default AWS managed SCP from the root level and all OU levels. For any specific exceptions, modify the SCP attached to that OU, and add the required AWS services to the allow list.
C. Use an SCP in Organizations to implement a deny list of AWS service. Apply this SCP at each OU level. Leave the default AWS managed SCP at the root level. For any specific executions for an OU, create a new SCP for that OU.
D. Use an SCP in Organizations to implement an allow list of AWS services. Apply this SCP at the root level. Remove the default AWS managed SCP from the root level and all OU levels. For any specific exceptions for an OU, modify the SCP attached to that OU, and add the required AWS services to the allow list.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 779

An organization has 4 people in the IT operations team who are responsible to manage the AWS infrastructure. The organization wants to setup that each user will have access to launch and manage an instance in a zone which the other user cannot modify.
Which of the below mentioned options is the best solution to set this up?

A. Create four AWS accounts and give each user access to a separate account.
B. Create an IAM user and allow them permission to launch an instance of a different sizes only.
C. Create four IAM users and four VPCs and allow each IAM user to have access to separate VPCs.
D. Create a VPC with four subnets and allow access to each subnet for the individual IAM user.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. The VPC also work with IAM and the organization can create IAM users who have access to various VPC services. The organization can setup access for the IAM user who can modify the security groups of the VPC.
The sample policy is given below:
{
"Version": "2012-10-17", "Statement":
[{ "Effect": "Allow",
"Action": "ec2:RunInstances", "Resource":
["arn:aws:ec2:region::image/ami-*", "arn:aws:ec2:region:account:subnet/subnet-1a2b3c4d", "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/sg-123abc123"] }]
} With this policy the user can create four subnets in separate zones and provide IAM user access to each subnet.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_IAM.html

QUESTION 780
The two policies that you attach to an IAM role are the access policy and the trust policy. The trust policy identifies who can assume the role and grants the permission in the AWS Lambda account principal by adding the _____ action.

A. aws:AssumeAdmin
B. lambda:InvokeAsync
C. sts:InvokeAsync
D. sts:AssumeRole

Correct Answer: D
Section: (none)

Explanation

Explanation/Reference:
Explanation:
The two policies that you attach to an IAM role are the access policy and the trust policy. Remember that adding an account to the trust policy of a role is only half of establishing the trust relationship. By default, no users in the trusted accounts can assume the role until the administrator for that account grants the users the permission to assume the role by adding the Amazon Resource Name (ARN) of the role to an Allow element for the sts:AssumeRole action.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_manage_modify.html

QUESTION 781
You are implementing a URL whitelisting system for a company that wants to restrict outbound HTTP'S connections tospecific domains from their EC2-hosted applications. You deploy a single EC2 instance running proxy software andconfigure It to accept traffic from all subnets and EC2 instances in the VPC. You configure the proxy to only pass throughtraffic to domains that you define in its whitelist configuration. You have a nightly maintenance window or 10 minutes whereall instances fetch new software updates. Each update Is about 200MB In size and there are 500 instances In the VPC thatroutinely fetch updates. After a few days you notice that some machines are failing to successfully download some, but notall of their updates within the maintenance window. The download URLs used for these updates are correctly listed in theproxy's whitelist configuration and you are able to access them manually using a web browser on the instances.
What might be happening? (Choose two.)

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 782
A large company with hundreds of AWS accounts has a newly established centralized internal process for purchasing new or modifying existing Reserved Instances. This process requires all business units that want to purchase or modify Reserved Instances to submit requests to a dedicated team for procurement or execution. Previously, business units would directly purchase or modify Reserved Instances in their own respective AWS accounts autonomously.
Which combination of steps should be taken to proactively enforce the new process in the MOST secure way possible?
(Choose two.)

A. Ensure all AWS accounts are part of an AWS Organizations structure operating in all features mode.
B. Use AWS Config to report on the attachment of an IAM policy that denies access to the ec2:PurchaseReservedInstancesOffering and ec2:ModifyReservedInstances actions.
C. In each AWS account, create an IAM policy with a DENY rule to the ec2:PurchaseReservedInstancesOffering and ec2:ModifyReservedInstances actions.
D. Create an SCP that contains a deny rule to the ec2:PurchaseReservedInstancesOffering and ec2:ModifyReservedInstances actions. Attach the SCP to each organizational unit (OU) of the AWS Organizations structure.
E. Ensure that all AWS accounts are part of an AWS Organizations structure operating in consolidated billing features mode.

Correct Answer: CE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 783
A company wants to replace its call center system with a solution built using AWS managed services. The company call center would like the solution to receive calls, create contact flows, and scale to handle growth projections. The call center would also like the solution to use deep learning capabilities to recognize the intent of the callers and handle basic tasks, reducing the need to speak to an agent. The solution should also be able to query business applications and provide relevant information back to callers as requested. Which services should the Solutions Architect use to build this solution?
(Choose three.)

A. Amazon Rekognition to identify who is calling.
B. Amazon Connect to create a cloud-based contact center.
C. Amazon Alexa for Business to build conversational interfaces.
D. AWS Lambda to integrate with internal systems.
E. Amazon Lex to recognize the intent of the caller.
F. Amazon SQS to add incoming callers to a queue.

Correct Answer: BDE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 784
How many cg1.4xlarge on-demand instances can a user run in one region without taking any limit increase approval from AWS?

A. 20
B. 2
C. 5
D. 10

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Generally, AWS EC2 allows running 20 on-demand instances and 100 spot instances at a time. This limit can be increased by requesting at https://aws.amazon.com/contact-us/ec2-request. Excluding certain types of instances, the limit is lower than mentioned above. For cg1.4xlarge, the user can run only 2 on-demand instances at a time.
Reference: http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#limits_ec2

QUESTION 785
A company plans to refactor a monolithic application into a modern application design deployed on AWS. The CI/CD pipeline needs to be upgraded to support the modern design for the application with the following requirements:
It should allow changes to be released several times every hour. It should be able to roll back the changes as quickly as possible.
Which design will meet these requirements?

A. Deploy a CI/CD pipeline that incorporates AMIs to contain the application and their configurations. Deploy the application by replacing Amazon EC2 instances.
B. Specify AWS Elastic Beanstalk to stage in a secondary environment as the deployment target for the CI/CD pipeline of the application. To deploy, swap the staging and production environment URLs.
C. Use AWS Systems Manager to re-provision the infrastructure for each deployment. Update the Amazon EC2 user data to pull the latest code artifact from Amazon S3 and use Amazon Route 53 weighted routing to point to the new environment.
D. Roll out the application updates as part of an Auto Scaling event using prebuilt AMIs. Use new versions of the AMIs to add instances, and phase out all instances that use the previous AMI version with the configured termination policy during a deployment event.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 786

A company is running a custom database in the AWS Cloud. The database uses Amazon EC2 for compute and uses Amazon Elastic Block Store (Amazon EBS) for storage. The database runs on the latest generation of EC2 instances and uses a General Purpose SSD (gp2) EBS volume for data.
The current data volume has the following characteristics:
The volume is 512 GB in size.
The volume never goes above 256 GB utilization.
The volume consistently uses around 1,500 IOPS.
A solutions architect needs to conduct an analysis of the current database storage layer and make a recommendation about ways to reduce cost.
Which solution will provide the MOST cost savings without impacting the performance of the database?

A. Convert the data volume to the Cloud HDD (sc1) type. Leave the volume as 512 GB. Set the volume IOPS to 1,500.
B. Convert the data volume to the Provisioned IOPS SSD (io2) type. Resize the volume to 256 GSet the volume IOPS to 1,500.
C. Convert the data volume to the Provisioned IOPS SSD (io2) Block Express type. Leave the volume as 512 GB. Set the volume IOPS to 1,500.
D. Convert the data volume to the General Purpose SSD (gp3) type. Resize the volume to 256 GB. Set the volume IOPS to 1,500.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 787
A company has a web application that uses Amazon API Gateway, AWS Lambda, and Amazon DynamoDB. A recent marketing campaign has increased demand. Monitoring software reports that many requests have significantly longer response times than before the marketing campaign.
A solutions architect enabled Amazon CloudWatch Logs for API Gateway and noticed that errors are occurring on 20% of the requests. In CloudWatch, the Lambda function Throttles metric represents 1% of the requests and the Errors metric represents 10% of the requests. Application logs indicate that, when errors occur, there is a call to DynamoDB.
What change should the solutions architect make to improve the current response times as the web application becomes more popular?

A. Increase the concurrency limit of the Lambda function
B. Implement DynamoDB auto scaling on the table
C. Increase the API Gateway throttle limit
D. Re-create the DynamoDB table with a better-partitioned primary index

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 788
You have a website which requires international presence and consequently you have set it up as follows.
It is hosted on 30 EC2 instances.
It is on in 15 regions around the globe. Each region has 2 instances. All the instances are a public hosted zone.
Which of the following is the best way to configure your site to maintain availability with minimum downtime if one of the 15 regions was to lose network connectivity for an extended period? (Choose two.)

A. Create a Route 53 Latency Based Routing Record set that resolves to an Elastic Load Balancer in each region and has the Evaluate Target Health flag set to true.
B. Create a Route 53 failover routing policy and configure an active-passive failover.
C. Create a Route 53 Failover Routing Policy and assign each resource record set a unique identifier and a relative weight.
D. Create a Route 53 Geolocation Routing Policy that resolves to an Elastic Load Balancer in each region and has the Evaluate Target Health flag set to false.

Correct Answer: AB
Section: (none)
Explanation

Explanation/Reference:
Explanation:
It is best to use the latency routing policy when you have resources in multiple Amazon EC2 data centers that perform the same function and you want Amazon Route 53 to respond to DNS queries with the resources that provide the best latency.
You could also use the failover routing policy (for public hosted zones only) when you want to configure an active-passive failover, in which one resource takes all traffic when it's available and the other resource takes all traffic when the first resource isn't available.
Reference: http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html#routing-policy-latency

QUESTION 789
A company that is new to AWS reports it has exhausted its service limits across several accounts that are on the Basic Support plan. The company would like to prevent this from happening in the future.
What is the MOST efficient way of monitoring and managing all service limits in the company's accounts?

A. Use Amazon CloudWatch and AWS Lambda to periodically calculate the limits across all linked accounts using AWS Trusted Advisor, provide notifications using Amazon SNS if the limits are close to exceeding the threshold.
B. Reach out to AWS Support to proactively increase the limits across all accounts. That way, the customer avoids creating and managing infrastructure just to raise the service limits.
C. Use Amazon CloudWatch and AWS Lambda to periodically calculate the limits across all linked accounts using AWS Trusted Advisor, programmatically increase the limits that are close to exceeding the threshold.
D. Use Amazon CloudWatch and AWS Lambda to periodically calculate the limits across all linked accounts using AWS Trusted Advisor, and use Amazon SNS for notifications if a limit is close to exceeding the threshold. Ensure that the accounts are using the AWS Business Support plan at a minimum.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 790
A company is building a software-as-a-service (SaaS) solution on AWS. The company has deployed an Amazon API Gateway REST API with AWS Lambda integration in multiple AWS Regions and in the same production account.
The company offers tiered pricing that gives customers the ability to pay for the capacity to make a certain number of API calls per second. The premium tier offers up to 3,000 calls per second, and customers are identified by a unique API key.
Several premium tier customers in various Regions report that they receive error responses of 429 Too Many Requests from multiple API methods during peak usage hours. Logs indicate that the Lambda function is never invoked.
What could be the cause of the error messages for these customers?

A. The Lambda function reached its concurrency limit.
B. The Lambda function its Region limit for concurrency.
C. The company reached its API Gateway account limit for calls per second.
D. The company reached its API Gateway default per-method limit for calls per second.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 791
A company is running its AWS infrastructure across two AWS Regions. The company has four VPCs in the eu-west-1 Region and has two VPCs in the us-east-1 Region. The company also has an onpremises data center in Europe that has two AWS Direct Connect connections in eu-west-1.
The company needs a solution in which Amazon EC2 instances in each VPC can connect to each other by using private IP addresses. Servers in the on-premises data center also must be able to connect to those VPCs by using private IP addresses.
What is the MOST cost-effective solution that meets these requirements?

A. Create an AWS Transit Gateway in each Region, and attach each VPC to the transit gateway in that Region. Create cross-Region peering between the transit gateways. Create two transit VIFs, and attach them to a single Direct Connect gateway. Associate each transit gateway with the Direct Connect gateway.
B. Create VPC peering between each VPC in the same Region. Create cross-Region peering between each VPC in different Regions. Create two private VIFs, and attach them to a single Direct Connect gateway. Associate each VPC with the Direct Connect gateway.

C. Create VPC peering between each VPC in the same Region. Create cross-Region peering between each VPC in different Regions. Create two public VIFs that are configured to route AWS IP addresses globally to on-premises servers.

D. Create an AWS Transit Gateway in each Region, and attach each VPC to the transit gateway in that Region. Create cross-Region peering between the transit gateways. Create two private VIFs, and attach them to a single Direct Connect gateway. Associate each VPC with the Direct Connect gateway.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 792
A company has multiple lines of business (LOBs) that roll up to the parent company. The company has asked its solutions architect to develop a solution with the following requirements:
Produce a single AWS invoice for all of the AWS accounts used by its LOBs.
The costs for each LOB account should be broken out on the invoice.
Provide the ability to restrict services and features in the LOB accounts, as defined by the company's governance policy.
Each LOB account should be delegated full administrator permissions, regardless of the governance policy.
Which combination of steps should the solutions architect take to meet these requirements? (Choose two.)

A. Use AWS Organizations to create an organization in the parent account for each LOB. Then, invite each LOB account to the appropriate organization.

B. Use AWS Organizations to create a single organization in the parent account. Then, invite each LOB's AWS account to pin the organization.

C. Implement service quotas to define the services and features that are permitted and apply the quotas to each LOB as appropriate.

D. Create an SCP that allows only approved services and features, then apply the policy to the LOB accounts. Enable consolidated billing in the parent account's billing console and link the LOB accounts.

Correct Answer: CD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 793
Amazon EC2 provides a repository of public data sets that can be seamlessly integrated into AWS cloud-based applications.
What is the monthly charge for using the public data sets?

A. A 1-time charge of 10$ for all the datasets.

B. 1$ per dataset per month
C. 10$ per month for all the datasets
D. There is no charge for using the public data sets

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 794
If a single condition within an IAM policy includes multiple values for one key, it will be evaluated using a logical_____.

A. OR
B. NAND
C. NOR
D. AND

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
If a single condition within an IAM policy includes multiple values for one key, it will be evaluated using a logical OR.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements.html

QUESTION 795
True or False: Amazon ElastiCache supports the Redis key-value store.

A. True, ElastiCache supports the Redis key-value store, but with limited functionalities.
B. False, ElastiCache does not support the Redis key-value store.
C. True, ElastiCache supports the Redis key-value store.
D. False, ElastiCache supports the Redis key-value store only if you are in a VPC environment.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
This is true. ElastiCache supports two open-source in-memory caching engines: 1. Memcached - a widely adopted memory object caching system. ElastiCache is protocol compliant with Memcached, so popular tools that you use today with existing Memcached environments will work seamlessly with the service. 2. Redis - a popular open-source in-memory key-value store that supports data structures such as sorted sets and lists. ElastiCache supports Master / Slave replication and Multi- AZ which can be used to achieve cross AZ redundancy.
Reference: https://aws.amazon.com/elasticache/

QUESTION 796
A company has created an account for individual Development teams, resulting in a total of 200 accounts. All accounts have a single virtual private cloud (VPC) in a single region with multiple microservices running in Docker containers that need to communicate with microservices in other accounts. The Security team requirements state that these microservices must not traverse the public internet, and only certain internal services should be allowed to call other individual services. If there is any denied network traffic for a service, the Security team must be notified of any denied requests, including the source IP.
How can connectivity be established between service while meeting the security requirements?

A. Create a VPC peering connection between the VPCs. Use security groups on the instances to allow traffic from the security group IDs that are permitted to call the microservice. Apply network ACLs and allow traffic from the local VPC and peered VPCs only. Within the task definition in Amazon ECS for each of the microservices, specify a log configuration by using the awslogs driver. Within Amazon CloudWatch Logs, create a metric filter and alarm off of the number of HTTP 403 responses. Create an alarm when the number of messages exceeds a threshold set by the Security team.
B. Ensure that no CIDR ranges are overlapping, and attach a virtual private gateway (VGW) to each VPC. Provision an IPsec tunnel between each VGW and enable route propagation on the route table. Configure security groups on each service to allow the CIDR ranges of the VPCs in the other accounts. Enable VPC Flow Logs, and use an Amazon CloudWatch Logs subscription filter for rejected traffic. Create an IAM role and allow the Security team to call the AssumeRole action for each account.
C. Deploy a transit VPC by using third-party marketplace VPN appliances running on Amazon EC2, dynamically routed VPN connections between the VPN appliance, and the virtual private gateways (VGWs) attached to each VPC within the region.
Adjust network ACLs to allow traffic from the local VPC only. Apply security groups to the microservices to allow traffic from the VPN appliances only. Install the awslogs agent on each VPN appliance, and configure logs to forward to Amazon CloudWatch Logs in the security account for the Security team to access.
D. Create a Network Load Balancer (NLB) for each microservice. Attach the NLB to a PrivateLink endpoint service and whitelist the accounts that will be consuming this service. Create an interface endpoint in the consumer VPC and associate a security group that allows only the security group IDs of the services authorized to call the producer service. On the producer services, create security groups for each microservice and allow only the CIDR range of the allowed services. Create VPC Flow Logs on each VPC to capture rejected traffic that will be delivered to an Amazon CloudWatch Logs group. Create a CloudWatch Logs subscription that streams the log data to a security account.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 797
Which of the following is true of Amazon EBS encryption keys?

A. Amazon EBS encryption uses the Customer Master Key (CMK) to create an AWS Key Management Service (AWS KMS) master key.
B. Amazon EBS encryption uses the EBS Magnetic key to create an AWS Key Management Service (AWS KMS) master key.
C. Amazon EBS encryption uses the EBS Magnetic key to create a Customer Master Key (CMK).
D. Amazon EBS encryption uses the AWS Key Management Service (AWS KMS) master key to create a Customer Master Key (CMK).

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Amazon EBS encryption uses AWS Key Management Service (AWS KMS) master keys when creating encrypted volumes and any snapshots created from your encrypted volumes.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonEBS.html

QUESTION 798
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 three.)

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: CDE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 799
A company recently deployed a new application that runs on a group of Amazon EC2 Linux instances in a VPC. In a peered VPC, the company launched an EC2

Linux instance that serves as a bastion host. The security group of the application instances allows access only on TCP port 22 from the private IP of the bastion host. The security group of the bastion host allows access to TCP port 22 from 0.0.0.0/0 so that system administrators can use SSH to remotely log in to the application instances from several branch offices.
While looking through operating system logs on the bastion host, a cloud engineer notices thousands of failed SSH logins to the bastion host from locations around the world. The cloud engineer wants to change how remote access is granted to the application instances and wants to meet the following requirements:
Eliminate brute-force SSH login attempts.
Retain a log of commands run during an SSH session. Retain the ability to forward ports.
Which solution meets these requirements for remote access to the application instances?

A. Configure the application instances to communicate with AWS Systems Manager. Grant access to the system administrators to use Session Manager to establish a session with the application instances. Terminate the bastion host.
B. Update the security group of the bastion host to allow traffic from only the public IP addresses of the branch offices.
C. Configure an AWS Client VPN endpoint and provision each system administrator with a certificate to establish a VPN connection to the application VPUpdate the security group of the application instances to allow traffic from only the Client VPN IPv4 CIDR. Terminate the bastion host.
D. Configure the application instances to communicate with AWS Systems Manager. Grant access to the system administrators to issue commands to the application instances by using Systems Manager Run Command. Terminate the bastion host.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 800
An organization has created 5 IAM users. The organization wants to give them the same login ID but different passwords.
How can the organization achieve this?

A. The organization should create each user in a separate region so that they have their own URL to login
B. The organization should create a separate login ID but give the IAM users the same alias so that each one can login with their alias
C. It is not possible to have the same login ID for multiple IAM users of the same account
D. The organization should create various groups and add each user with the same login ID to different groups. The user can login with their own group ID

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. Whenever the organization is creating an IAM user, there should be a unique ID for each user. It is not possible to have the same login ID for multiple users. The names of users, groups, roles, instance profiles must be alphanumeric, including the following common characters: plus (+), equal (=), comma (,), period (.), at (@), and dash (-).
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_SettingUpUser.html

QUESTION 801
Cognito Sync is an AWS service that you can use to synchronize user profile data across mobile devices without requiring your own backend. When the device is online, you can synchronize data.
If you also set up push sync, what does it allow you to do?

A. Notify other devices that a user profile is available across multiple devices
B. Synchronize user profile data with less latency
C. Notify other devices immediately that an update is available
D. Synchronize online data faster

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Cognito Sync is an AWS service that you can use to synchronize user profile data across mobile devices without requiring your own backend. When the device is online, you can synchronize data, and if you have also set up push sync, notify other devices immediately that an update is available.
Reference: http://docs.aws.amazon.com/cognito/devguide/sync/

QUESTION 802
A company hosts a web application on AWS in the us-east-1 Region. The application servers are distributed across three Availability Zones behind an Application Load Balancer. The database is hosted in MySQL database on an Amazon EC2 instance. A solutions architect needs to design a cross-Region data recovery solution using AWS services with an RTO of less than 5 minutes and an RPO of less than 1 minute. The solutions architect is deploying application servers in us-west-2, and has configured Amazon Route 53 health checks and DNS failover to us-west-2.
Which additional step should the solutions architect take?

A. Migrate the database to an Amazon RDS for MySQL instance with a cross-Region read replica in us-west-2.
B. Migrate the database to an Amazon Aurora global database with the primary in us-east-1 and the secondary in us-west-2.
C. Migrate the database to an Amazon RDS for MySQL instance with a Multi-AZ deployment.
D. Create a MySQL standby database on an Amazon EC2 instance in us-west-2.

Correct Answer: B

Section: (none)
Explanation

Explanation/Reference:

QUESTION 803
A Solutions Architect must update an application environment within AWS Elastic Beanstalk using a blue/green deployment methodology. The Solutions Architect creates an environment that is identical to the existing application environment and deploys the application to the new environment. What should be done next to complete the update?

A. Redirect to the new environment using Amazon Route 53
B. Select the Swap Environment URLs option
C. Replace the Auto Scaling launch configuration
D. Update the DNS records to point to the green environment

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.CNAMESwap.html

QUESTION 804
A company is migrating an application to AWS. It wants to use fully managed services as much as possible during the migration. The company needs to store large, important documents within the application with the following requirements:
The data must be highly durable and available.
The data must always be encrypted at rest and in transit.
The encryption key must be managed by the company and rotated periodically.
Which of the following solutions should the Solutions Architect recommend?

A. Deploy the storage gateway to AWS in file gateway mode. Use Amazon EBS volume encryption using an AWS KMS key to encrypt the storage gateway volumes.
B. Use Amazon S3 with a bucket policy to enforce HTTPS for connections to the bucket and to enforce server-sideencryption and AWS KMS for object encryption.
C. Use Amazon DynamoDB with SSL to connect to DynamoDB. Use an AWS KMS key to encrypt DynamoDB objects at rest.
D. Deploy instances with Amazon EBS volumes attached to store this data. Use EBS volume encryption using an AWS KMS key to encrypt the data.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 805
A solutions architect has been assigned to migrate a 50 TB Oracle data warehouse that contains sales data from onpremises to Amazon Redshift. Major updates to the sales data occur on the final calendar day of the month. For the remainder of the month, the data warehouse only receives minor daily updates and is primarily used for reading and reporting. Because of this, the migration process must start on the first day of the month and must be complete before the next set of updates occur. This provides approximately 30 days to complete the migration and ensure that the minor daily changes have been synchronized with the Amazon Redshift data warehouse. Because the migration cannot impact normal business network operations, the bandwidth allocated to the migration for moving data over the internet is 50 Mbps. The company wants to keep data migration costs low.
Which steps will allow the solutions architect to perform the migration within the specified timeline?

A. Install Oracle database software on an Amazon EC2 instance. Configure VPN connectivity between AWS and the company's data center. Configure the Oracle database running on Amazon EC2 to join the Oracle Real Application Clusters (RAC). When the Oracle database on Amazon EC2 finishes synchronizing, create an AWS DMS ongoing replication task to migrate the data from the Oracle database on Amazon EC2 to Amazon Redshift. Verify the data migration is complete and perform the cut over to Amazon Redshift.
B. Create an AWS Snowball import job. Export a backup of the Oracle data warehouse. Copy the exported data to the Snowball device. Return the Snowball device to AWS. Create an Amazon RDS for Oracle database and restore the backup file to that RDS instance. Create an AWS DMS task to migrate the data from the RDS for Oracle database to Amazon Redshift. Copy daily incremental backups from Oracle in the data center to the RDS for Oracle database over the internet.
Verify the data migration is complete and perform the cut over to Amazon Redshift.
C. Install Oracle database software on an Amazon EC2 instance. To minimize the migration time, configure VPN connectivity between AWS and the company's data center by provisioning a 1 Gbps AWS Direct Connect connection. Configure the Oracle database running on Amazon EC2 to be a read replica of the data center Oracle database. Start the synchronization process between the company's on-premises data center and the Oracle database on Amazon EC2. When the Oracle database on Amazon EC2 is synchronized with the on-premises database, create an AWS DMS ongoing replication task to migrate the data from the Oracle database read replica that is running on Amazon EC2 to Amazon Redshift. Verify the data migration is complete and perform the cut over to Amazon Redshift.
D. Create an AWS Snowball import job. Configure a server in the company's data center with an extraction agent. Use AWS SCT to manage the extraction agent and convert the Oracle schema to an Amazon Redshift schema. Create a new project in AWS SCT using the registered data extraction agent. Create a local task and an AWS DMS task in AWS SCT with replication of ongoing changes. Copy data to the Snowball device and return the Snowball device to AWS. Allow AWS DMS to copy data from Amazon S3 to Amazon Redshift. Verify that the data migration is complete and perform the cut over to Amazon Redshift.

Correct Answer: A
Section: (none)

Explanation

Explanation/Reference:

QUESTION 806
Attempts, one of the three types of items associated with the schedule pipeline in the AWS Data Pipeline, provides robust data management.
Which of the following statements is NOT true about Attempts?

A. Attempts provide robust data management.
B. AWS Data Pipeline retries a failed operation until the count of retries reaches the maximum number of allowed retry attempts.
C. An AWS Data Pipeline Attempt object compiles the pipeline components to create a set of actionable instances.
D. AWS Data Pipeline Attempt objects track the various attempts, results, and failure reasons if applicable.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Attempts, one of the three types of items associated with a schedule pipeline in AWS Data Pipeline, provides robust data management. AWS Data Pipeline retries a failed operation. It continues to do so until the task reaches the maximum number of allowed retry attempts. Attempt objects track the various attempts, results, and failure reasons if applicable. Essentially, it is the instance with a counter. AWS Data Pipeline performs retries using the same resources from the previous attempts, such as Amazon EMR clusters and EC2 instances.
Reference: http://docs.aws.amazon.com/datapipeline/latest/DeveloperGuide/dp-how-tasks-scheduled.html

QUESTION 807
By default, temporary security credentials for an IAM user are valid for a maximum of 12 hours, but you can request a duration as long as _____ hours.

A. 24
B. 36
C. 10
D. 48

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

Explanation:
By default, temporary security credentials for an IAM user are valid for a maximum of 12 hours, but you can request a duration as short as 15 minutes or as long as 36 hours.
Reference: http://docs.aws.amazon.com/STS/latest/UsingSTS/CreatingSessionTokens.html

QUESTION 808
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 809
A company has multiple AWS accounts hosting IT applications. An Amazon CloudWatch Logs agent is installed on all Amazon EC2 instances. The company wants to aggregate all security events in a centralized AWS account dedicated to log storage.
Security Administrators need to perform near-real-time gathering and correlating of events across multiple AWS accounts.
Which solution satisfies these requirements?

A. Create a Log Audit IAM role in each application AWS account with permissions to view CloudWatch Logs, configure an AWS Lambda function to assume the Log Audit role, and perform an hourly export of CloudWatch Logs data to an Amazon S3 bucket in the logging AWS account.
B. Configure CloudWatch Logs streams in each application AWS account to forward events to CloudWatch Logs in the logging AWS account. In the logging AWS account, subscribe an Amazon Kinesis Data Firehose stream to Amazon CloudWatch Events, and use the stream to persist log data in Amazon S3.
C. Create Amazon Kinesis Data Streams in the logging account, subscribe the stream to CloudWatch Logs streams in each application AWS account, configure

an Amazon Kinesis Data Firehose delivery stream with the Data Streams as its source, and persist the log data in an Amazon S3 bucket inside the logging AWS account.

D. Configure CloudWatch Logs agents to publish data to an Amazon Kinesis Data Firehose stream in the logging AWS account, use an AWS Lambda function to read messages from the stream and push messages to Data Firehose, and persist the data in Amazon S3.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://noise.getoto.net/2018/03/03/central-logging-in-multi-account-environments/

QUESTION 810
True or False: "In the context of Amazon ElastiCache, from the application's point of view, connecting to the cluster configuration endpoint is no different than connecting directly to an individual cache node."

A. True, from the application's point of view, connecting to the cluster configuration endpoint is no different than connecting directly to an individual cache node since, each has a unique node identifier.

B. True, from the application's point of view, connecting to the cluster configuration endpoint is no different than connecting directly to an individual cache node.

C. False, you can connect to a cache node, but not to a cluster configuration endpoint.

D. False, you can connect to a cluster configuration endpoint, but not to a cache node.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
This is true. From the application's point of view, connecting to the cluster configuration endpoint is no different than connecting directly to an individual cache node. In the process of connecting to cache nodes, the application resolves the configuration endpoint's DNS name. Because the configuration endpoint maintains CNAME entries for all of the cache nodes, the DNS name resolves to one of the nodes; the client can then connect to that node.
Reference: http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/AutoDiscovery.HowAutoDiscoveryWorks.html

QUESTION 811
A company has an application that generates a weather forecast that is updated every 15 minutes with an output resolution of 1 billion unique positions, each approximately 20 bytes in size (20 Gigabytes per forecast). Every hour, the forecast data is globally accessed approximately 5 million times (1,400 requests per second), and up to 10 times more during weather events. The forecast data is overwritten every update. Users of the current weather forecast application expect responses to queries to be returned in less than two seconds for each request.

Which design meets the required request rate and response time?

A. Store forecast locations in an Amazon ES cluster. Use an Amazon CloudFront distribution targeting an Amazon API Gateway endpoint with AWS Lambda functions responding to queries as the origin. Enable API caching on the API Gateway stage with a cache-control timeout set for 15 minutes.
B. Store forecast locations in an Amazon EFS volume. Create an Amazon CloudFront distribution that targets an Elastic Load Balancing group of an Auto Scaling fleet of Amazon EC2 instances that have mounted the Amazon EFS volume. Set the cache-control timeout for 15 minutes in the CloudFront distribution.
C. Store forecast locations in an Amazon ES cluster. Use an Amazon CloudFront distribution targeting an API Gateway endpoint with AWS Lambda functions responding to queries as the origin. Create an Amazon Lambda@Edge function that caches the data locally at edge locations for 15 minutes.
D. Store forecast locations in Amazon S3 as individual objects. Create an Amazon CloudFront distribution targeting an Elastic Load Balancing group of an Auto Scaling fleet of EC2 instances, querying the origin of the S3 object. Set the cachecontrol timeout for 15 minutes in the CloudFront distribution.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://aws.amazon.com/blogs/networking-and-content-delivery/lambdaedge-design-best-practices/

QUESTION 812
True or False: The Amazon ElastiCache clusters are not available for use in VPC at this time.

A. TRUE
B. True, but they are available only in the GovCloud.
C. True, but they are available only on request
D. FALSE

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Amazon Elasticache clusters can be run in an Amazon VPC. With Amazon VPC, you can define a virtual network topology and customize the network configuration to closely resemble a traditional network that you might operate in your own datacenter. You can now take advantage of the manageability, availability and scalability benefits of Amazon ElastiCache Clusters in your own isolated network. The same functionality of Amazon ElastiCache, including automatic failure detection, recovery, scaling, auto discovery, Amazon CloudWatch metrics, and software patching, are now available in Amazon VPC.

Reference: http://aws.amazon.com/about-aws/whats-new/2012/12/20/amazon-elasticache-announces-support-for-amazon-vpc/

QUESTION 813
A company's CISO has asked a Solutions Architect to re-engineer the company's current CI/CD practices to make sure patch deployments to its applications can happen as quickly as possible with minimal downtime if vulnerabilities are discovered. The company must also be able to quickly roll back a change in case of errors.
The web application is deployed in a fleet of Amazon EC2 instances behind an Application Load Balancer. The company is currently using GitHub to host the application source code, and has configured an AWS CodeBuild project to build the application. The company also intends to use AWS CodePipeline to trigger builds from GitHub commits using the existing CodeBuild project.
What CI/CD configuration meets all of the requirements?

A. Configure CodePipeline with a deploy stage using AWS CodeDeploy configured for in-place deployment. Monitor the newly deployed code, and, if there are any issues, push another code update.
B. Configure CodePipeline with a deploy stage using AWS CodeDeploy configured for blue/green deployments. Monitor the newly deployed code, and, if there are any issues, trigger a manual rollback using CodeDeploy.
C. Configure CodePipeline with a deploy stage using AWS CloudFormation to create a pipeline for test and production stacks. Monitor the newly deployed code, and, if there are any issues, push another code update.
D. Configure the CodePipeline with a deploy stage using AWS OpsWorks and in-place deployments. Monitor the newly deployed code, and, if there are any issues, push another code update.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 814
A user has configured EBS volume with PIOPS. The user is not experiencing the optimal throughput.
Which of the following could not be factor affecting I/O performance of that EBS volume?

A. EBS bandwidth of dedicated instance exceeding the PIOPS
B. EBS volume size
C. EC2 bandwidth
D. Instance type is not EBS optimized

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
If the user is not experiencing the expected IOPS or throughput that is provisioned, ensure that the EC2 bandwidth is not the limiting factor, the instance is EBS-optimized (or include 10 Gigabit network connectivity) and the instance type EBS dedicated bandwidth exceeds the IOPS more than he has provisioned.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-io-characteristics.html

QUESTION 815
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:

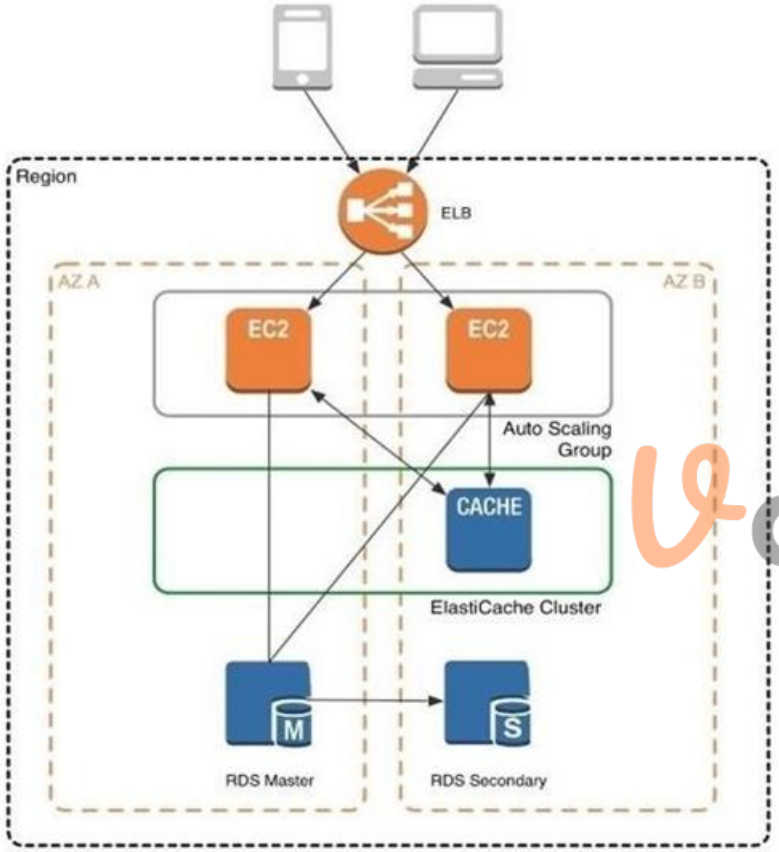

In this architecture diagram, the Amazon EC2 application instances are in an Auto Scaling group, located behind a load balancer using Elastic Load Balancing, which distributes requests among the instances. As requests come into a given EC2 instance, that EC2 instance is responsible for communicating 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.

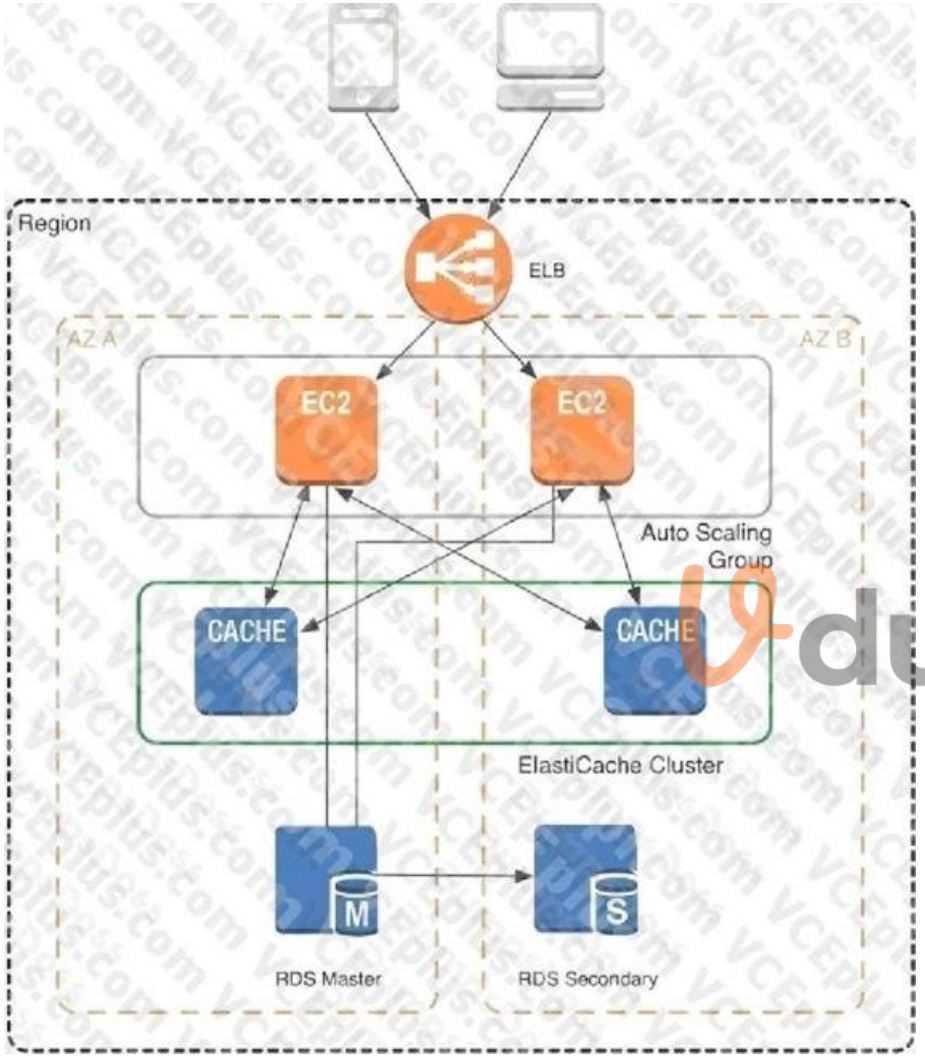

When you launch an ElastiCache cluster, you can choose the Availability Zone(s) that the cluster lives in. For best performance, you should configure your cluster to use the same Availability Zones as your application servers. To launch an ElastiCache cluster in a specific Availability Zone, make sure 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 tradeoff 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:

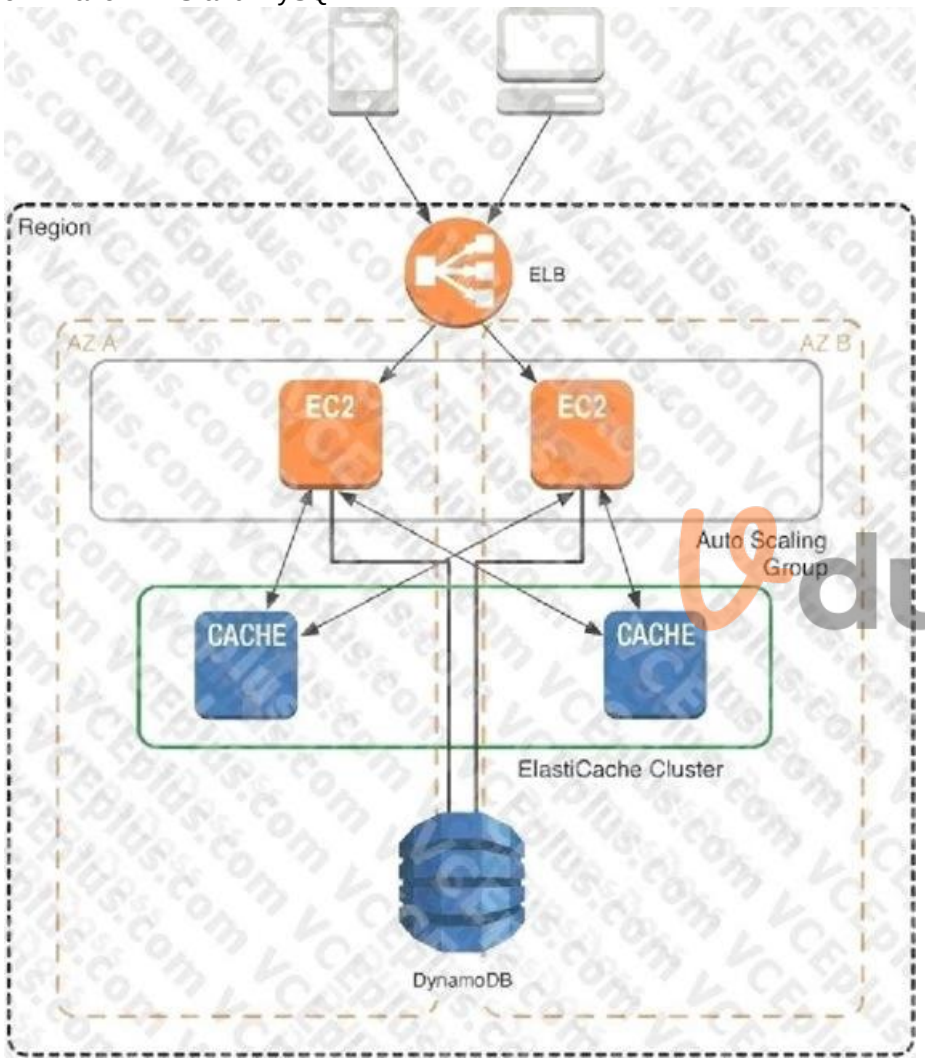

This combination of DynamoDB and ElastiCache is very popular with mobile and game companies, because DynamoDB allows for higher write throughput at lower cost than traditional relational databases.

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 816
The following are AWS Storage services? (Choose two.)

A. AWS Relational Database Service (AWS RDS)
B. AWS ElastiCache
C. AWS Glacier
D. AWS Import/Export

Correct Answer: CD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 817
A user is running a batch process on EBS backed EC2 instances. The batch process launches few EC2 instances to process Hadoop Map reduce jobs which can run between 50 ?600 minutes or sometimes for even more time. The user wants a configuration that can terminate the instance only when the process is completed.
How can the user configure this with CloudWatch?

A. Configure a job which terminates all instances after 600 minutes
B. It is not possible to terminate instances automatically
C. Configure the CloudWatch action to terminate the instance when the CPU utilization falls below 5%
D. Set up the CloudWatch with Auto Scaling to terminate all the instances

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Amazon CloudWatch alarm watches a single metric over a time period that the user specifies and performs one or more actions based on the value of the metric relative to a given threshold over a number of time periods. The user can setup an action which terminates the instances when their CPU utilization is below a certain threshold for a certain period of time. The EC2 action can either terminate or stop the instance as part of the EC2 action.
Reference: http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/UsingAlarmActions.html

QUESTION 818

Is it possible to load data from Amazon DynamoDB into Amazon Redshift?

A. No, you cannot load all the data from DynamoDB table to a Redshift table as it limited by size constraints.
B. No
C. No, DynamoDB data types do not correspond directly with those of Amazon Redshift.
D. Yes

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Yes. When you copy data from an Amazon DynamoDB table into Amazon Redshift, you can perform complex data analysis queries on that data. This includes joins with other tables in your Amazon Redshift cluster.
Reference: http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/RedshiftforDynamoDB.html

QUESTION 819
A company would like to implement a serverless application by using Amazon API Gateway, AWS Lambda, and Amazon DynamoDB. They deployed a proof of concept and stated that the average response time is greater than what their upstream services can accept. Amazon CloudWatch metrics did not indicate any issues with DynamoDB but showed that some Lambda functions were hitting their timeout.
Which of the following actions should the Solutions Architect consider to improve performance? (Choose two.)

A. Configure the AWS Lambda function to reuse containers to avoid unnecessary startup time.
B. Increase the amount of memory and adjust the timeout on the Lambda function. Complete performance testing to identify the ideal memory and timeout configuration for the Lambda function.
C. Create an Amazon ElastiCache cluster running Memcached, and configure the Lambda function for VPC integration with access to the Amazon ElastiCache cluster.
D. Enable API cache on the appropriate stage in Amazon API Gateway, and override the TTL for individual methods that require a lower TTL than the entire stage.
E. Increase the amount of CPU, and adjust the timeout on the Lambda function. Complete performance testing to identify the ideal CPU and timeout configuration for the Lambda function.

Correct Answer: BD
Section: (none)
Explanation

Explanation/Reference:
Explanation:

Reference:
https://lumigo.io/blog/aws-lambda-timeout-best-practices/

QUESTION 820

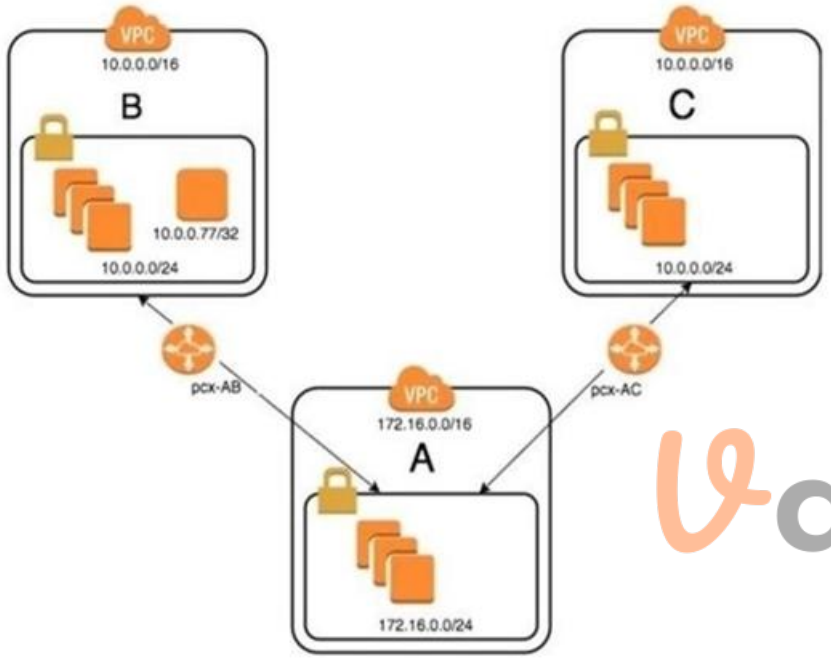

An organization has recently grown through acquisitions. Two of the purchased companies use the same IP CIDR range.
There is a new short-term requirement to allow AnyCompany A (VPC-A) to communicate with a server that has the IP address 10.0.0.77 in AnyCompany B (VPC-B). AnyCompany A must also communicate with all resources in AnyCompany C (VPC-C). The Network team has created the VPC peer links, but it is having issues with communications between VPC-A and VPC-B. After an investigation, the team believes that the routing tables in the VPCs are incorrect.
What configuration will allow AnyCompany A to communicate with AnyCompany C in addition to the database in AnyCompany B?

A. On VPC-A, create a static route for the VPC-B CIDR range (10.0.0.0/24) across VPC peer pcx-AB.
Create a static route of 10.0.0.0/16 across VPC peer pcx-AC.
On VPC-B, create a static route for VPC-A CIDR (172.16.0.0/24) on peer pcx-AB.
On VPC-C, create a static route for VPC-A CIDR (172.16.0.0/24) across peer pcx-AC.

B. On VPC-A, enable dynamic route propagation on pcx-AB and pcx-AC.
On VPC-B, enable dynamic route propagation and use security groups to allow only the IP address 10.0.0.77/32 on VPC peer pcx-AB. On VPC-C, enable dynamic route propagation with VPC-A on peer pcx-AC.

C. On VPC-A, create network access control lists that block the IP address 10.0.0.77/32 on VPC peer pcx-AC.

On VPC-A, create a static route for VPC-B CIDR (10.0.0.0/24) on pcx-AB and a static route for VPC-C CIDR (10.0.0.0/24) on pcx-AC.
On VPC-B, create a static route for VPC-A CIDR (172.16.0.0/24) on peer pcx-AB.
On VPC-C, create a static route for VPC-A CIDR (172.16.0.0/24) across peer pcx-AC.

D. On VPC-A, create a static route for the VPC-B (10.0.0.77/32) database across VPC peer pcx-AB.
Create a static route for the VPC-C CIDR on VPC peer pcx-AC.
On VPC-B, create a static route for VPC-A CIDR (172.16.0.0/24) on peer pcx-AB.
On VPC-C, create a static route for VPC-A CIDR (172.16.0.0/24) across peer pcx-AC.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 821
In the context of IAM roles for Amazon EC2, which of the following NOT true about delegating permission to make API requests?

A. You cannot create an IAM role.
B. You can have the application retrieve a set of temporary credentials and use them.
C. You can specify the role when you launch your instances.
D. You can define which accounts or AWS services can assume the role.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Amazon designed IAM roles so that your applications can securely make API requests from your instances, without requiring you to manage the security credentials that the applications use.
Instead of creating and distributing your AWS credentials, you can delegate permission to make API requests using IAM roles as follows: Create an IAM role. Define which accounts or AWS services can assume the role. Define which API actions and resources the application can use after assuming the role. Specify the role when you launch your instances. Have the application retrieve a set of temporary credentials and use them.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html

QUESTION 822
A company is migrating mobile banking applications to run on Amazon EC2 instances in a VPC. Backend service applications run in an on-premises data center. The data center has an AWS Direct Connect connection into AWS. The applications that run in the VPC need to resolve DNS requests to an on-premises Active Directory domain that runs in the data center.

Which solution will meet these requirements with the LEAST administrative overhead?

A. Provision a set of EC2 instances across two Availability Zones in the VPC as caching DNS servers to resolve DNS queries from the application servers within the VPC.
B. Provision an Amazon Route 53 private hosted zone. Configure NS records that point to on-premises DNS servers.
C. Create DNS endpoints by using Amazon Route 53 Resolver Add conditional forwarding rules to resolve DNS namespaces between the on-premises data center and the VPC.
D. Provision a new Active Directory domain controller in the VPC with a bidirectional trust between this new domain and the on-premises Active Directory domain.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/blogs/security/how-to-set-up-dns-resolution-between-on-premises-networks-and-awsusing-aws-directory-service-and-amazon-route-53/

QUESTION 823
A company has an application that generates reports and stores them in an Amazon bucket Amazon S3 bucket. When a user accesses their report, the application generates a signed URL to allow the user to download the report. The company's security team has discovered that the files are public and that anyone can download them without authentication. The company has suspended the generation of new reports until the problem is resolved.
Which set of action will immediately remediate the security issue without impacting the application's normal workflow?

A. Create an AWS Lambda function that applies all policy for users who are not authenticated. Create a scheduled event to invoke the Lambda function.
B. Review the AWS Trusted advisor bucket permissions check and implement the recommend actions.
C. Run a script that puts a Private ACL on all of the object in the bucket.
D. Use the Block Public Access feature in Amazon S3 to set the IgnorePublicAcis option to TRUE on the bucket.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 824
A user is planning to host a web server as well as an app server on a single EC2 instance which is a part of the public subnet of a VPC.

How can the user setup to have two separate public IPs and separate security groups for both the application as well as the web server?

A. Launch VPC with two separate subnets and make the instance a part of both the subnets.
B. Launch a VPC instance with two network interfaces. Assign a separate security group and elastic IP to them.
C. Launch a VPC instance with two network interfaces. Assign a separate security group to each and AWS will assign a separate public IP to them.
D. Launch a VPC with ELB such that it redirects requests to separate VPC instances of the public subnet.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
If you need to host multiple websites (with different IPs) on a single EC2 instance, the following is the suggested method from AWS. Launch a VPC instance with two network interfaces.
Assign elastic IPs from VPC EIP pool to those interfaces (Because, when the user has attached more than one network interface with an instance, AWS cannot assign public IPs to them.) Assign separate Security Groups if separate Security Groups are needed This scenario also helps for operating network appliances, such as firewalls or load balancers that have multiple private IP addresses for each network interface.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/MultipleIP.html

QUESTION 825
A company has deployed an application to multiple environments in AWS, including production and testing. The company has separate accounts for production and testing, and users are allowed to create additional application users for team members or services, as needed. The Security team has asked the Operations team for better isolation between production and testing with centralized controls on security credentials and improved management of permissions between environments.
Which of the following options would MOST securely accomplish this goal?

A. Create a new AWS account to hold user and service accounts, such as an identity account. Create users and groups in the identity account. Create roles with appropriate permissions in the production and testing accounts. Add the identity account to the trust policies for the roles.
B. Modify permissions in the production and testing accounts to limit creating new IAM users to members of the Operations team. Set a strong IAM password policy on each account. Create new IAM users and groups in each account to limit developer access to just the services required to complete their job function.
C. Create a script that runs on each account that checks user accounts for adherence to a security policy. Disable any user or service accounts that do not comply.
D. Create all user accounts in the production account. Create roles for access in the production account and testing accounts. Grant cross-account access from the production account to the testing account.

Correct Answer: A
Section: (none)

Explanation

Explanation/Reference:
Explanation:
Reference:
https://aws.amazon.com/ru/blogs/security/how-to-centralize-and-automate-iam-policy-creation-in-sandbox-development-andtest-environments/

QUESTION 826
A company has more than 100 AWS accounts, with one VPC per account, that need outbound HTTPS connectivity to theinternet. The current design contains one NAT gateway per Availability Zone (AZ) in each VPC. To reduce costs and obtaininformation about outbound traffic, management has asked for a new architecture for internet access.
Which solution will meet the current needs, and continue to grow as new accounts are provisioned, while reducing costs?

A. Create a transit VPC across two AZs using a third-party routing appliance. Create a VPN connection to each VPC. Default route internet traffic to the transit VPC.
B. Create multiple hosted-private AWS Direct Connect VIFs, one per account, each with a Direct Connect gateway. Default route internet traffic back to an on-premises router to route to the internet.
C. Create a central VPC for outbound internet traffic. Use VPC peering to default route to a set of redundant NAT gateway in the central VPC.
D. Create a proxy fleet in a central VPC account. Create an AWS PrivateLink endpoint service in the central VPC. Use PrivateLink interface for internet connectivity through the proxy fleet.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference:
https://aws.amazon.com/blogs/networking-and-content-delivery/how-to-use-aws-privatelink-to-secure-and-scale-webfiltering-using-explicit-proxy/

QUESTION 827
A large company recently experienced an unexpected increase in Amazon RDS and Amazon DynamoDB costs. The company needs to increase visibility into details of AWS Billing and Cost Management. There are various accounts associated with AWS Organizations, including many development and production accounts. There is no consistent tagging strategy across the organization, but there are guidelines in place that require all infrastructure to be deployed using AWS CloudFormation with consistent tagging. Management requires cost center numbers and project ID numbers for all existing and future DynamoDB tables and RDS instances.
Which strategy should the solutions architect provide to meet these requirements?

A. Use Tag Editor to tag existing resources. Create cost allocation tags to define the cost center and project ID and allow 24 hours for tags to propagate to existing resources.

B. Use an AWS Config rule to alert the finance team of untagged resources. Create a centralized AWS Lambda based solution to tag untagged RDS databases and DynamoDB resources every hour using a cross-account role.
C. Use Tag Editor to tag existing resources. Create cost allocation tags to define the cost center and project ID. Use SCPs to restrict resource creation that do not have the cost center and project ID on the resource.
D. Create cost allocation tags to define the cost center and project ID and allow 24 hours for tags to propagate to existing resources. Update existing federated roles to restrict privileges to provision resources that do not include the cost center and project ID on the resource.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/cost-alloc-tags.html

QUESTION 828
A company's lease of a colocated storage facility will expire in 90 days. The company wants to move to AWS to avoid signing a contract extension. The company's environment consists of 200 virtual machines and a NAS with 40 TB of data.
Most of the data is archival, yet instant access is required when data is requested. Leadership wants to ensure minimal downtime during the migration. Each virtual machine has a number of customized configurations. The company's existing 1 Gbps network connection is mostly idle, especially after business hours.
Which combination of steps should the company take to migrate to AWS while minimizing downtime and operational impact? (Choose two.)

A. Use new Amazon EC2 instances and reinstall all application code.
B. Use AWS SMS to migrate the virtual machines.
C. Use AWS Storage Gateway to migrate the data to cloud-native storage.
D. Use AWS Snowball to migrate the data.
E. Use AWS SMS to copy the infrequently accessed data from the NAS.

Correct Answer: AD
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/snowball/latest/ug/transfer-data.html

QUESTION 829
A company has 50 AWS accounts that are members of an organization in AWS Organizations. Each account contains multiple VPCs. The company wants to use AWS Transit Gateway to establish connectivity between the VPCs in each member account. Each time a new member account is created, the company wants to

automate the process of creating a new VPC and a transit gateway attachment.
Which combination of steps will meet these requirements? (Choose two.)

A. From the management account, share the transit gateway with member accounts by using AWS Resource Access Manager.
B. From the management account, share the transit gateway with member accounts by using an AWS Organizations SCP.
C. Launch an AWS CloudFormation stack set from the management account that automatically creates a new VPC and a VPC transit gateway attachment in a member account. Associate the attachment with the transit gateway in the management account by using the transit gateway ID.
D. Launch an AWS CloudFormation stack set from the management account that automatically creates a new VPC and a peering transit gateway attachment in a member account. Share the attachment with the transit gateway in the management account by using a transit gateway service-linked role.
E. From the management account, share the transit gateway with member accounts by using AWS Service Catalog.

Correct Answer: AC
Section: (none)
Explanation

Explanation/Reference:

QUESTION 830
The Security team needs to provide a team of interns with an AWS environment so they can build a serverless video transcoding application. The project will use Amazon S3, AWS Lambda, Amazon API Gateway, Amazon Cognito, Amazon DynamoDB, and Amazon Elastic Transcoder.
The interns should be able to create and configure the necessary resources, but they may not have access to create or modify AWS IAM roles. The Solutions Architect creates a policy and attaches it to the interns' group.
How should the Security team configure the environment to ensure that the interns are self-sufficient?

A. Create a policy that allows creation of project-related resources only. Create roles with required service permissions, which are assumable by the services.
B. Create a policy that allows creation of all project-related resources, including roles that allow access only to specified resources.
C. Create roles with the required service permissions, which are assumable by the services. Have the interns create and use a bastion host to create the project resources in the project subnet only.
D. Create a policy that allows creation of project-related resources only. Require the interns to raise a request for roles to be created with the Security team. The interns will provide the requirements for the permissions to be set in the role.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 831
A weather service provides high-resolution weather maps from a web application hosted on AWS in the eu-west-1 Region.
The weather maps are updated frequently and stored in Amazon S3 along with static HTML content. The web application is fronted by Amazon CloudFront.
The company recently expanded to serve users in the us-east-1 Region, and these new users report that viewing their respective weather maps is slow from time to time.
Which combination of steps will resolve the us-east-1 performance issues? (Choose two.)

A. Configure the AWS Global Accelerator endpoint for the S3 bucket in eu-west-1. Configure endpoint groups for TCP ports 80 and 443 in us-east-1.
B. Create a new S3 bucket in us-east-1. Configure S3 cross-Region replication to synchronize from the S3 bucket in euwest- 1.
C. Use Lambda@Edge to modify requests from North America to use the S3 Transfer Acceleration endpoint in us-east-1.
D. Use Lambda@Edge to modify requests from North America to use the S3 bucket in us-east-1.
E. Configure the AWS Global Accelerator endpoint for us-east-1 as an origin on the CloudFront distribution. Use Lambda@Edge to modify requests from North America to use the new origin.

Correct Answer: BC
Section: (none)
Explanation

Explanation/Reference:

QUESTION 832
A solutions architect is evaluating the reliability of a recently migrated application running on AWS. The front end is hosted on Amazon S3 and accelerated by Amazon CloudFront. The application layer is running in a stateless Docker container on an Amazon EC2 On-Demand Instance with an Elastic IP address. The storage layer is a MongoDB database running on an EC2 Reserved Instance in the same Availability Zone as the application layer.
Which combination of steps should the solutions architect take to eliminate single points of failure with minimal application code changes? (Choose two.)

A. Create a REST API in Amazon API Gateway and use AWS Lambda functions as the application layer
B. Create an Application Load Balancer and migrate the Docker container to AWS Fargate
C. Migrate the storage layer to Amazon DynamoDB
D. Migrate the storage layer to Amazon DocumentDB (with MongoDB compatibility)
E. Create an Application Load Balancer and move the storage layer to an EC2 Auto Scaling group

Correct Answer: AE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 833
You are designing a personal document-archiving solution for your global enterprise with thousands of employees. 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 834
A company is using AWS Organizations to manage multiple AWS accounts. For security purposes, the company requires the creation of an Amazon Simple Notification Service (Amazon SNS) topic that enables integration with a third-party alerting system in all the Organizations member accounts.
A solutions architect used an AWS CloudFormation template to create the SNS topic and stack sets to automate the deployment of CloudFormation stacks. Trusted access has been enabled in Organizations.
What should the solutions architect do to deploy the CloudFormation StackSets in all AWS accounts?

A. Create a stack set in the Organizations member accounts. Use service-managed permissions. Set deployment options to deploy to an organization. Use CloudFormation StackSets drift detection.
B. Create stacks in the Organizations member accounts. Use self-service permissions. Set deployment options to deploy to an organization. Enable the CloudFormation StackSets automatic deployment.
C. Create a stack set in the Organizations master account. Use service-managed permissions. Set deployment options to deploy to the organization. Enable CloudFormation StackSets automatic deployment.
D. Create stacks in the Organizations master account. Use service-managed permissions. Set deployment options to deploy to the organization. Enable CloudFormation StackSets drift detection.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-orgs-manage-autodeployment.html

QUESTION 835
In the context of AWS CloudFormation, which of the following statements is correct?

A. Actual resource names are a combination of the resource ID, stack, and logical resource name.
B. Actual resource name is the stack resource name.
C. Actual resource name is the logical resource name.
D. Actual resource names are a combination of the stack and logical resource name.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In AWS CloudFormation, actual resource names are a combination of the stack and logical resource name. This allows multiple stacks to be created from a template without fear of name collisions between AWS resources.
Reference: https://aws.amazon.com/cloudformation/faqs/

QUESTION 836
You have deployed a three-tier web application in a VPC with a CIDR 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 two.)

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:
Explanation:
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

QUESTION 837
A company is using AWS CloudFormation as its deployment tool for all application. It stages all application binaries and templates within Amazon S3 bucket with versioning enabled. Developers have access to an Amazon EC2 instance that hosts the integrated development (IDE). The Developers download the application binaries from Amazon S3 to the EC2 instance, make changes, and upload the binaries to an S3 bucket after running the unit tests locally. The developers want to improve the existing deployment mechanism and implement CI/CD using AWS CodePipeline.
The developers have the following requirements:
Use AWS CodeCommit for source control.
Automate unit testing and security scanning.
Alert the Developers when unit tests fail.
Turn application features on and off, and customize deployment dynamically as part of CI/CD. Have the lead Developer provide approval before deploying an application.
Which solution will meet these requirements?

A. Use AWS CodeBuild to run tests and security scans. Use an Amazon EventBridge rule to send Amazon SNS alerts to the Developers when unit tests fail. Write AWS Cloud Developer kit (AWS CDK) constructs for different solution features, and use a manifest file to turn features on and off in the AWS CDK application. Use a manual approval stage in the pipeline to allow the lead Developer to approve applications.
B. Use AWS Lambda to run unit tests and security scans. Use Lambda in a subsequent stage in the pipeline to send Amazon SNS alerts to the developers when unit tests fail. Write AWS Amplify plugins for different solution features and utilize user prompts to turn features on and off. Use Amazon SES in the pipeline to allow the lead developer to approve applications.
C. Use Jenkins to run unit tests and security scans. Use an Amazon EventBridge rule in the pipeline to send Amazon SES alerts to the developers when unit tests fail. Use AWS CloudFormation nested stacks for different solution features and parameters to turn features on and off. Use AWS Lambda in the pipeline to allow the lead developer to approve applications.
D. Use AWS CodeDeploy to run unit tests and security scans. Use an Amazon CloudWatch alarm in the pipeline to send Amazon SNS alerts to the developers when unit tests fail. Use Docker images for different solution features and the AWS CLI to turn features on and off. Use a manual approval stage in the pipeline to allow the lead developer to approve applications.

Correct Answer: C

Section: (none)
Explanation

Explanation/Reference:

QUESTION 838
A company is planning to migrate an existing high performance computing (HPC) solution to the AWS Cloud. The existing solution consists of a 12-node cluster running Linux with high speed interconnectivity developed on a single rack. A solutions architect needs to optimize the performance of the HPC cluster.
Which combination of steps will meet these requirements? (Choose two.)

A. Deploy instances across at least three Availability Zones.
B. Deploy Amazon EC2 instances in a placement group.
C. Use Amazon EC2 instances that support Elastic Fabric Adapter (EFA).
D. Use Amazon EC2 instances that support burstable performance.
E. Enable CPU hyperthreading.

Correct Answer: BE
Section: (none)
Explanation

Explanation/Reference:

QUESTION 839
The CISO of a large enterprise with multiple IT departments, each with its own AWS account, wants one central place where AWS permissions for users can be managed and users authentication credentials can be synchronized with the company's existing on-premises solution.
Which solution will meet the CISO's requirements?

A. Define AWS IAM roles based on the functional responsibilities of the users in a central account. Create a SAML-based identity management provider. Map users in the on-premises groups to IAM roles. Establish trust relationships between the other accounts and the central account.
B. Deploy a common set of AWS IAM users, groups, roles, and policies in all of the AWS accounts using AWS Organizations. Implement federation between the on-premises identity provider and the AWS accounts.
C. Use AWS Organizations in a centralized account to define service control policies (SCPs). Create a SAML-based identity management provider in each account and map users in the on-premises groups to AWS IAM roles.
D. Perform a thorough analysis of the user base and create AWS IAM users accounts that have the necessary permissions. Set up a process to provision and deprovision accounts based on data in the on-premises solution.

Correct Answer: C

Section: (none)
Explanation

Explanation/Reference:

QUESTION 840
A user has created a VPC with CIDR 20.0.0.0/16 using the VPC wizard. The user has created public and VPN only subnets along with hardware VPN access to connect to the user's data center. The user has not yet launched any instance as well as modified or deleted any setup. He wants to delete this VPC from the console.
Will the console allow the user to delete the VPC?

A. Yes, the user can detach the virtual private gateway and then use the VPC console to delete the VPC.
B. No, since the NAT instance is running, the user cannot delete the VPC.
C. Yes, the user can use the CLI to delete the VPC that will detach the virtual private gateway automatically.
D. No, the VPC console needs to be accessed using an administrator account to delete the VPC.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
You can delete your VPC at any time (for example, if you decide it's too small). However, you must terminate all instances in the VPC first. When you delete a VPC using the VPC console, Amazon deletes all its components, such as subnets, security groups, network ACLs, route tables, Internet gateways, VPC peering connections, and DHCP options. If you have a VPN connection, you don't have to delete it or the other components related to the VPN (such as the customer gateway and virtual private gateway).
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html#VPC_Deleting

QUESTION 841
An enterprise company wants to implement cost controls for all its accounts in AWS Organizations, which has full features enabled. The company has mapped organizational units (OUs) to its business units, and it wants to bill these business units for their individual AWS spending. There has been a recent spike in the company's AWS bill, which is generating attention from the Finance team. A Solutions Architect needs to investigate the cause of the spike while designing a solution that will track AWS costs in Organizations and generate a notification to the required teams if costs from a business unit exceed a specific monetary threshold.
Which solution will meet these requirements?

A. Use Cost Explorer to troubleshoot the reason for the additional costs. Set up an AWS Lambda function to monitor the company's AWS bill by each AWS account in an OU. Store the threshold amount set by the Finance team in the AWS Systems Manager Parameter Store. Write the custom rules in the Lambda function to verify any hidden costs for the AWS accounts. Trigger a notification from the Lambda function to an Amazon SNS topic when a budget threshold is

breached.

B. Use AWS Trusted Advisor to troubleshoot the reason for the additional costs. Set up an AWS Lambda function to monitor the company's AWS bill by each AWS account in an OU. Store the threshold amount set by the Finance team in the AWS Systems Manager Parameter Store. Write custom rules in the Lambda function to verify any hidden costs for the AWS accounts. Trigger an email to the required teams from the Lambda function using Amazon SNS when a budget threshold is breached.

C. Use Cost Explorer to troubleshoot the reason for the additional costs. Create a budget using AWS Budgets with the monetary amount set by the Finance team for each OU by grouping the linked accounts. Configure an Amazon SNS notification to the required teams in the budget.

D. Use AWS Trusted Advisor to troubleshoot the reason for the additional costs. Create a budget using AWS Budgets with the monetary amount set by the Finance team for each OU by grouping the linked accounts. Add the Amazon EC2 instance types to be used in the company as a budget filter. Configure an Amazon SNS topic with a subscription for the Finance team email address to receive budget notifications.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 842
Which of the following is the Amazon Resource Name (ARN) condition operator that can be used within an Identity and Access Management (IAM) policy to check the case-insensitive matching of the ARN?

A. ArnCheck

B. ArnMatch

C. ArnCase

D. ArnLike

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Amazon Resource Name (ARN) condition operators let you construct Condition elements that restrict access based on comparing a key to an ARN. ArnLike, for instance, is a case-insensitive matching of the ARN. Each of the six colon-delimited components of the ARN is checked separately and each can include a multi-character match wildcard (*) or a singlecharacter match wildcard (?).
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_ElementDescriptions.html

QUESTION 843

A company hosts a large on-premises MySQL database at its main office that supports an issue tracking system used by employees around the world. The company already uses AWS for some workloads and has created an Amazon Route 53 entry for the database endpoint that points to the on-premises database. Management is concerned about the database being a single point of failure and wants a solutions architect to migrate the database to AWS without any data loss or downtime.
Which set of actions should the solutions architect implement?

A. Create an Amazon Aurora DB cluster. Use AWS Database Migration Service (AWS DMS) to do a full load from the onpremises database to Aurora. Update the Route 53 entry for the database to point to the Aurora cluster endpoint, and shut down the on-premises database.
B. During nonbusiness hours, shut down the on-premises database and create a backup. Restore this backup to an Amazon Aurora DB cluster. When the restoration is complete, update the Route 53 entry for the database to point to the Aurora cluster endpoint, and shut down the on-premises database.
C. Create an Amazon Aurora DB cluster. Use AWS Database Migration Service (AWS DMS) to do a full load with continuous replication from the on-premises database to Aurora. When the migration is complete, update the Route 53 entry for the database to point to the Aurora cluster endpoint, and shut down the on-premises database.
D. Create a backup of the database and restore it to an Amazon Aurora multi-master cluster. This Aurora cluster will be in a master-master replication configuration with the on-premises database. Update the Route 53 entry for the database to point to the Aurora cluster endpoint, and shut down the on-premises database.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 844
A user has launched two EBS backed EC2 instances in the US-East-1a region. The user wants to change the zone of one of the instances. How can the user change it?

A. It is not possible to change the zone of an instance after it is launched
B. From the AWS EC2 console, select the Actions - > Change zones and specify the new zone
C. The zone can only be modified using the AWS CLI
D. Stop one of the instances and change the availability zone

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:

With AWS EC2, when a user is launching an instance he can select the availability zone (AZ) at the time of launch. If the zone is not selected, AWS selects it on behalf of the user. Once the instance is launched, the user cannot change the zone of that instance unless he creates an AMI of that instance and launches a new instance from it.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html

QUESTION 845
A company has registered 10 new domain names. The company uses the domains for online marketing. The company needs a solution that will redirect online visitors to a specific URL for each domain. All domains and target URLS are defined in a JSON document. All DNS records are managed by Amazon Route 53.
A solutions architect must implement a redirect service that accepts HTTP and HTTPS requests.
Which combination of steps should the solutions architect take to meet these requirements with the LEAST amount of operational effort? (Choose three.)

A. Create a dynamic webpage that runs on an Amazon EC2 instance. Configure the webpage to use the JSON document in combination with the event message to look up and respond with a redirect URL.
B. Create an Application Load Balancer that includes HTTP and HTTPS listeners.
C. Create an AWS Lambda function that uses the JSON document in combination with the event message to look up and respond with a redirect URL.
D. Use an Amazon API Gateway API with a custom domain to publish an AWS Lambda function.
E. Create an Amazon CloudFront distribution. Deploy a Lambda@Edge function.
F. Create an SSL certificate by using AWS Certificate Manager (ACM). Include the domains as Subject Alternative Names.

Correct Answer: ABF
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/premiumsupport/knowledge-center/route-53-redirect-to-another-domain/

QUESTION 846
A company has a large on-premises Apache Hadoop cluster with a 20 PB HDFS database. The cluster is growing every quarter by roughly 200 instances and 1 PB. The company's goals are to enable resiliency for its Hadoop data, limit the impact of losing cluster nodes, and significantly reduce costs. The current cluster runs 24/7 and supports a variety of analysis workloads, including interactive queries and batch processing.
Which solution would meet these requirements with the LEAST expense and down time?

A. Use AWS Snowmobile to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster initially sized to handle the interactive workload based on historical data from the onpremises cluster. Store the data on EMRFS. Minimize costs using Reserved Instances for master and core nodes and Spot Instances for task nodes, and auto scale task nodes based on Amazon CloudWatch metrics. Create job-specific, optimized clusters for batch workloads that are similarly optimized.
B. Use AWS Snowmobile to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster of a similar size and configuration to the current cluster. Store the data on EMRFS. Minimize costs by using Reserved Instances.
As the workload grows each quarter, purchase additional Reserved Instances and add to the cluster.

C. Use AWS Snowball to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster initially sized to handle the interactive workloads based on historical data from the onpremises cluster. Store the data on EMRFS.
Minimize costs using Reserved Instances for master and core nodes and Spot Instances for task nodes, and auto scale task nodes based on Amazon CloudWatch metrics. Create job-specific, optimized clusters for batch workloads that are similarly optimized.

D. Use AWS Direct Connect to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster initially sized to handle the interactive workload based on historical data from the on-premises cluster. Store the data on EMRFS. Minimize costs using Reserved Instances for master and core nodes and Spot Instances for task nodes, and auto scale task nodes based on Amazon CloudWatch metrics. Create job-specific, optimized clusters for batch workloads that are similarly optimized.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
To migrate large datasets of 10 PB or more in a single location, you should use Snowmobile. For datasets less than 10 PB or distributed in multiple locations, you should use Snowball. In addition, you should evaluate the amount of available bandwidth in your network backbone. If you have a high speed backbone with hundreds of Gb/s of spare throughput, then you can use Snowmobile to migrate the large datasets all at once. If you have limited bandwidth on your backbone, you should consider using multiple Snowballs to migrate the data incrementally.

QUESTION 847
In regard to DynamoDB, when you create a table with a hash-and-range key.

A. You must define one or more Local secondary indexes on that table
B. You must define one or more Global secondary indexes on that table
C. You can optionally define one or more secondary indexes on that table
D. You must define one or more secondary indexes on that table

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
When you create a table with a hash-and-range key, you can optionally define one or more secondary indexes on that table.
A secondary index lets you query the data in the table using an alternate key, in addition to queries against the primary key.
Reference: http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DataModel.html

QUESTION 848
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 three.)

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 849
A North American company with headquarters on the East Coast is deploying a new web application running on Amazon EC2 in the us-east-1 Region. The application should dynamically scale to meet user demand and maintain resiliency.
Additionally, the application must have disaster recover capabilities in an active-passive configuration with the us-west-1 Region.
Which steps should a solutions architect take after creating a VPC in the us-east-1 Region?

A. Create a VPC in the us-west-1 Region. Use inter-Region VPC peering to connect both VPCs. Deploy an Application Load Balancer (ALB) spanning multiple Availability Zones (AZs) to the VPC in the useast-1 Region. Deploy EC2 instances across multiple AZs in each Region as part of an Auto Scaling group spanning both VPCs and served by the ALB.
B. Deploy an Application Load Balancer (ALB) spanning multiple Availability Zones (AZs) to the VPC in the us-east-1 Region. Deploy EC2 instances across multiple AZs as part of an Auto Scaling group served by the ALDeploy the same solution to the us-west-1 Region. Create an Amazon Route 53 record set with a failover routing policy and health checks enabled to provide high availability across both Regions.
C. Create a VPC in the us-west-1 Region. Use inter-Region VPC peering to connect both VPCs. Deploy an Application Load Balancer (ALB) that spans both VPCs. Deploy EC2 instances across multiple Availability Zones as part of an Auto Scaling group in each VPC served by the ALB. Create an Amazon Route 53 record that points to the ALB.
D. Deploy an Application Load Balancer (ALB) spanning multiple Availability Zones (AZs) to the VPC in the us-east-1 Region. Deploy EC2 instances across multiple AZs as part of an Auto Scaling group served by the ALB. Deploy the same solution to the us-west-1 Region. Create separate Amazon Route 53 records in each Region that point to the ALB in the Region. Use Route 53 health checks to provide high availability across both Regions.

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

QUESTION 850
What is the name of licensing model in which I can use your existing Oracle Database licenses to run Oracle deployments on Amazon RDS?

A. Bring Your Own License
B. Role Bases License
C. Enterprise License
D. License Included

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/oracle/

QUESTION 851
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 usergenerated 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 852
To get started using AWS Direct Connect, in which of the following steps do you configure Border Gateway Protocol (BGP)?

A. Complete the Cross Connect
B. Configure Redundant Connections with AWS Direct Connect
C. Create a Virtual Interface
D. Download Router Configuration

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In AWS Direct Connect, your network must support Border Gateway Protocol (BGP) and BGP MD5 authentication, and you need to provide a private Autonomous System Number (ASN) for that to connect to Amazon Virtual Private Cloud (VPC). To connect to public AWS products such as Amazon EC2 and Amazon S3, you will also need to provide a public ASN that you own (preferred) or a private ASN. You have to configure BGP in the Create a Virtual Interface step.
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/getstarted.html#createvirtualinterface

QUESTION 853
A retail company is operating its ecommerce application on AWS. The application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The company uses an Amazon RDS DB instance as the database backend. Amazon CloudFront is configured with one origin that points to the ALB. Static content is cached. Amazon Route 53 is used to host all public zones.
After an update of the application, the ALB occasionally returns a 502 status code (Bad Gateway) error. The root cause is malformed HTTP headers that are returned to the ALB. The webpage returns successfully when a solutions architect reloads the webpage immediately after the error occurs.
While the company is working on the problem, the solutions architect needs to provide a custom error page instead of the standard ALB error page to visitors.
Which combination of steps will meet this requirement with the LEAST amount of operational overhead? (Choose two.)

A. Create an Amazon S3 bucket. Configure the S3 bucket to host a static webpage. Upload the custom error pages to Amazon S3.

B. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function if the ALB health check response Target FailedHealthChecks is greater than 0. Configure the Lambda function to modify the forwarding rule at the ALB to point to a publicly accessible web server.
C. Modify the existing Amazon Route 53 records by adding health checks. Configure a fallback target if the health check fails. Modify DNS records to point to a publicly accessible webpage.
D. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function if the ALB health check response Elb.InternalError is greater than 0. Configure the Lambda function to modify the forwarding rule at the ALB to point to a public accessible web server.
E. Add a custom error response by configuring a CloudFront custom error page. Modify DNS records to point to a publicly accessible web page.

Correct Answer: DE
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/blogs/compute/scaling-amazon-ecs-services-automatically-using-amazon-cloudwatchand-aws-lambda/

QUESTION 854
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 855
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 three.)

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: BCD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 856
A Solutions Architect is designing a system that will collect and store data from 2,000 internet-connected sensors. Each sensor produces 1 KB of data every second. The data must be available for analysis within a few seconds of it being sent to the system and stored for analysis indefinitely. Which is the MOST cost-effective solution for collecting and storing the data?

A. Put each record in Amazon Kinesis Data Streams. Use an AWS Lambda function to write each record to an object in Amazon S3 with a prefix that organizes the records by hour and hashes the record's key. Analyze recent data from Kinesis Data Streams and historical data from Amazon S3.
B. Put each record in Amazon Kinesis Data Streams. Set up Amazon Kinesis Data Firehouse to read records from the stream and group them into objects in Amazon S3. Analyze recent data from Kinesis Data Streams and historical data from Amazon S3.
C. Put each record into an Amazon DynamoDB table. Analyze the recent data by querying the table. Use an AWS Lambda function connected to a DynamoDB stream to group records together, write them into objects in Amazon S3, and then delete the record from the DynamoDB table. Analyze recent data from the DynamoDB table and historical data from Amazon S3
D. Put each record into an object in Amazon S3 with a prefix what organizes the records by hour and hashes the record's key. Use S3 lifecycle management to transition objects to S3 infrequent access storage to reduce storage costs. Analyze recent and historical data by accessing the data in Amazon S3

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 857
A company maintains a restaurant review website. The website is a single-page application where files are stored in Amazon S3 and delivered using Amazon CloudFront. The company receives several fake postings every day that are manually removed.
The security team has identified that most of the fake posts are from bots with IP addresses that have a bad reputation within the same global region. The team

needs to create a solution to help restrict the bots from accessing the website.
Which strategy should a solutions architect use?

A. Use AWS Firewall Manage to control the CloudFront distribution security settings. Create a geographical block rule and associate it with Firewall Manager.
B. Associate an AWS WAF web ACL with the CloudFront distribution. Select the managed Amazon IP reputation rule group for the web ACL with a deny action.
C. Use AWS Firewall Manager to control the CloudFront distribution security settings. Select the managed Amazon IP reputation rule group and associate it with Firewall Manager with a deny action.
D. Associate an AWS WAF web ACL with the CloudFront distribution. Create a rule group for the web ACL with a geographical match statement with a deny action.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 858
In DynamoDB, a projection is_____.

A. systematic transformation of the latitudes and longitudes of the locations inside your table
B. importing data from your file to a table
C. exporting data from a table to your file
D. the set of attributes that is copied from a table into a secondary index

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
In DynamoDB, a projection is the set of attributes that is copied from a table into a secondary index.
Reference: http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GSI.html

QUESTION 859
In DynamoDB, which of the following operations is not possible by the console?

A. Updating an item
B. Copying an item

C. Blocking an item
D. Deleting an item

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
By using the console to manage DynamoDB, you can perform the following: adding an item, deleting an item, updating an item, and copying an item.
Reference: http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AddUpdateDeleteItems.html

QUESTION 860
A company wants to ensure that the workloads for each of its business units have complete autonomy and a minimal blast radius in AWS. The Security team must be able to control access to the resources and services in the account to ensure that particular services are not used by the business units.
How can a Solutions Architect achieve the isolation requirements?

A. Create individual accounts for each business unit and add the account to an OU in AWS Organizations. Modify the OU to ensure that the particular services are blocked. Federate each account with an IdP, and create separate roles for the business units and the Security team.
B. Create individual accounts for each business unit. Federate each account with an IdP and create separate roles and policies for business units and the Security team.
C. Create one shared account for the entire company. Create separate VPCs for each business unit. Create individual IAM policies and resource tags for each business unit. Federate each account with an IdP, and create separate roles for the business units and the Security team.
D. Create one shared account for the entire company. Create individual IAM policies and resource tags for each business unit. Federate the account with an IdP, and create separate roles for the business units and the Security team.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:

QUESTION 861
A company is in the process of implementing AWS Organizations to constrain its developers to use only Amazon EC2, Amazon S3, and Amazon DynamoDB. The Developers account resides in a dedicated organizational unit (OU). The Solutions Architect has implemented the following SCP on the Developers account:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowEC2",
            "Effect": "Allow",
            "Action": "ec2:*",
            "Resource": "*"
        },
        {
            "Sid": "AllowDynamoDB",
            "Effect": "Allow",
            "Action": "dynamodb:*",
            "Resource": "*"
        },
        {
            "Sid": "AllowS3",
            "Effect": "Allow",
            "Action": "s3:*",
            "Resource": "*"
        }
    ]
}
```

When this policy is deployed, IAM users in the Developers account are still able to use AWS services that are not listed in the policy.
What should the Solutions Architect do to eliminate the Developers' ability to use services outside the scope of this policy?

A. Create an explicit deny statement for each AWS service that should be constrained.
B. Remove the FullAWSAccess SCP from the Developer account's OU.
C. Modify the FullAWSAccess SCP to explicitly deny all services.
D. Add an explicit deny statement using a wildcard to the end of the SCP.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 862
A company runs a software-as-a-service (SaaS) application on AWS. The application consists of AWS Lambda functions and an Amazon RDS for MySQL Multi-AZ database. During market events, the application has a much higher workload than normal. Users notice slow response times during the peak periods because of many database connections. The company needs to improve the scalable performance and availability of the database.
Which solution meets these requirements?

A. Create an Amazon CloudWatch alarm action that triggers a Lambda function to add an Amazon RDS for MySQL read replica when resource utilization hits a threshold.
B. Migrate the database to Amazon Aurora, and add a read replica. Add a database connection pool outside of the Lambda handler function.
C. Migrate the database to Amazon Aurora, and add a read replica. Use Amazon Route 53 weighted records.
D. Migrate the database to Amazon Aurora, and add an Aurora Replica. Configure Amazon RDS Proxy to manage database connection pools.

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Reference: https://aws.amazon.com/blogs/database/tag/aws-lambda/feed/

QUESTION 863
How does AWS Data Pipeline execute activities on on-premise resources or AWS resources that you manage?

A. By supplying a Task Runner package that can be installed on your on-premise hosts
B. None of these
C. By supplying a Task Runner file that the resources can access for execution
D. By supplying a Task Runner json script that can be installed on your on-premise hosts

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
To enable running activities using on-premise resources, AWS Data Pipeline does the following: It supply a Task Runner package that can be installed on your on-premise hosts. This package continuously polls the AWS Data Pipeline service for work to perform. When it's time to run a particular activity on your on-premise resources, it will issue the appropriate command to the Task Runner.
Reference:
https://aws.amazon.com/datapipeline/faqs/

QUESTION 864
A Provisioned IOPS volume must be at least _____ GB in size:

A. 20
B. 10

C. 50
D. 1

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A Provisioned IOPS volume must be at least 10 GB in size
Reference: http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/Storage.html

QUESTION 865
A company is planning to deploy a new business analytics application that requires 10,000 hours of compute time each month. The compute resources can have flexible availability, but must be as costeffective as possible. The company will also provide a reporting service to distribute analytics reports, which needs to run at all times.
How should the Solutions Architect design a solution that meets these requirements?

A. Deploy the reporting service on a Spot Fleet. Deploy the analytics application as a container in Amazon ECS with AWS Fargate as the compute option. Set the analytics application to use a custom metric with Service Auto Scaling.
B. Deploy the reporting service on an On-Demand Instance. Deploy the analytics application as a container in AWS Batch with AWS Fargate as the compute option. Set the analytics application to use a custom metric with Service Auto Scaling.
C. Deploy the reporting service as a container in Amazon ECS with AWS Fargate as the compute option. Deploy the analytics application on a Spot Fleet. Set the analytics application to use a custom metric with Amazon EC2 Auto Scaling applied to the Spot Fleet.
D. Deploy the reporting service as a container in Amazon ECS with AWS Fargate as the compute option. Deploy the analytics application on an On-Demand Instance and purchase a Reserved Instance with a 3-year term. Set the analytics application to use a custom metric with Amazon EC2 Auto Scaling applied to the On-Demand Instance.

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:

QUESTION 866
A Solutions Architect is migrating a 10 TB PostgreSQL database to Amazon RDS for PostgreSQL. The company's internet link is 50 MB with a VPN in the Amazon VPC, and the Solutions Architect needs to migrate the data and synchronize the changes before the cutover. The cutover must take place within an 8-day period. What is the LEAST complex method of migrating the database securely and reliably?

A. Order an AWS Snowball device and copy the database using the AWS DMS. When the database is available in Amazon S3, use AWS DMS to load it to Amazon RDS, and configure a job to synchronize changes before the cutover.
B. Create an AWS DMS job to continuously replicate the data from on premises to AWS. Cutover to Amazon RDS after the data is synchronized.
C. Order an AWS Snowball device and copy a database dump to the device. After the data has been copied to Amazon S3, import it to the Amazon RDS instance. Set up log shipping over a VPN to synchronize changes before the cutover.
D. Order an AWS Snowball device and copy the database by using the AWS Schema Conversion Tool. When the data is available in Amazon S3, use AWS DMS to load it to Amazon RDS, and configure a job to synchronize changes before the cutover.

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

QUESTION 867
A company is using multiple AWS accounts. The company has a shared service account and several other accounts for different projects.
A team has a VPC in a project account. The team wants to connect this VPC to a corporate network through an AWS Direct Connect gateway that exists in the shared services account. The team wants to automatically perform a virtual private gateway association with the Direct Connect gateway by using an already-tested AWS Lambda function while deploying its VPC networking stack. The Lambda function code can assume a role by using AWS Security Token Service (AWS STS).
The team is using AWS CloudFormation to deploy its infrastructure.
Which combination of steps will meet these requirements? (Choose three.)

A. Deploy the Lambda function to the project account. Update the Lambda function's IAM role with the directconnect:* permission.
B. Create a cross-account IAM role in the shared services account that grants the Lambda function the directconnect:* permission. Add the sts:AssumeRole permission to the IAM role that is associated with the Lambda function in the shared services account.
C. Add a custom resource to the CloudFormation networking stack that references the Lambda function in the project account.
D. Deploy the Lambda function that is performing the association to the shared services account. Update the Lambda function's IAM role with the directconnect:* permission.
E. Create a cross-account IAM role in the shared services account that grants the sts:AssumeRole permission to the Lambda function with the directconnect:* permission acting as a resource. Add the sts:AssumeRole permission with this cross-account IAM role as a resource to the IAM role that belongs to the Lambda function in the project account.
F. Add a custom resource to the CloudFormation networking stack that references the Lambda function in the shared services account.

Correct Answer: CEF
Section: (none)
Explanation

Explanation/Reference:

QUESTION 868
Which of the following are characteristics of Amazon VPC subnets? (Choose two.)

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: BD
Section: (none)
Explanation

Explanation/Reference:

QUESTION 869
Which of following IAM policy elements lets you specify an exception to a list of actions?

A. NotException
B. ExceptionAction
C. Exception
D. NotAction

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:
Explanation:
The NotAction element lets you specify an exception to a list of actions.