Amazon-Web-Services

Exam Questions SAP-C02

AWS Certified Solutions Architect - Professional

NEW QUESTION 1
- (Exam Topic 1)
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 as possible?

A. Deploy a Lambda@Edge function to sort parameters by name and force them to be lowercas
B. Select the CloudFront viewer request trigger to invoke the function.
C. Update the CloudFront distribution to disable caching based on query string parameters.
D. 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.
E. Update the CloudFront distribution to specify casing-insensitive query string processing.

Answer: A

Explanation:
https://docs.amazonaws.cn/en_us/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html#lambda-ex Before CloudFront serves content from the cache it will trigger any Lambda function associated with the Viewer Request, in which we can normalize parameters.
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html#lambda-examp

NEW QUESTION 2
- (Exam Topic 1)
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 cost
B. Convert the Oracle database into a single Amazon RDS reserved DB instance.
C. Create an Auto Scaling group for the front end with a combination of On-Demand and Spot Instances to reduce cost
D. Create two additional copies of the database instance, then distribute the databases in separate AZs.
E. Create an Auto Scaling group for the front end with a combination of On-Demand and Spot Instances to reduce cost
F. Convert the tables in the Oracle database into Amazon DynamoDB tables.
G. Convert the On-Demand Instances into Spot Instances to reduce costs for the front en
H. Convert the tables in the Oracle database into Amazon DynamoDB tables.

Answer: C

Explanation:
Combination of On-Demand and Spot Instances + DynamoDB.

NEW QUESTION 3
- (Exam Topic 1)
A company wants to deploy an AWS WAF solution to manage AWS WAF rules across multiple AWS accounts. The accounts are managed under different OUs in AWS Organizations.
Administrators must be able to add or remove accounts or OUs from managed AWS WAF rule sets as needed Administrators also must have the ability to automatically update and remediate noncompliant AWS WAF rules in all accounts
Which solution meets these requirements with the LEAST amount of operational overhead?

A. Use AWS Firewall Manager to manage AWS WAF rules across accounts in the organizatio
B. Use an AWS Systems Manager Parameter Store parameter to store account numbers and OUs to manage Update the parameter as needed to add or remove accounts or OUs Use an Amazon EventBridge (Amazon CloudWatch Events) rule to identify any changes to the parameter and to invoke an AWS Lambda function to update the security policy in the Firewall Manager administrative account
C. Deploy an organization-wide AWS Config rule that requires all resources in the selected OUs to associate the AWS WAF rule
D. Deploy automated remediation actions by using AWS Lambda to fix noncompliant resources Deploy AWS WAF rules by using an AWS CloudFormation stack set to target the same OUs where the AWS Config rule is applied.
E. Create AWS WAF rules in the management account of the organization Use AWS Lambda environment variables to store account numbers and OUs to manage Update environment variables as needed to add or remove accounts or OUs Create cross-account IAM roles in member accounts Assume the rotes by using AWS Security Token Service (AWS STS) in the Lambda function to create and update AWS WAF rules in the member accounts.
F. Use AWS Control Tower to manage AWS WAF rules across accounts in the organization Use AWS Key Management Service (AWS KMS) to store account numbers and OUs to manage Update AWS KMS as needed to add or remove accounts or OUs Create IAM users in member accounts Allow AWS Control Tower in the management account to use the access key and secret access key to create and update AWS WAF rules in the member accounts

Answer: B

NEW QUESTION 4
- (Exam Topic 1)
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 lo 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 owne
B. Add each business unit to an Amazon SNS topic for each aler
C. Use Cost Explorer in each account to create monthly reports for each business unit.
D. Configure AWS Budgets in the organization's master account and configure budget alerts that are grouped by application, environment, and owne
E. Add each business unit to an Amazon SNS topic for each aler
F. Use Cost Explorer in the organization's master account to create monthly reports for each business unit.
G. Configure AWS Budgets in each account and configure budget alerts lhat are grouped by application, environment, and owne
H. Add each business unit to an Amazon SNS topic for each aler
I. Use the AWS Billing and Cost Management dashboard in each account to create monthly reports for each business unit.
J. Enable AWS Cost and Usage Reports in the organization's master account and configure reports grouped by application, environment, and owne
K. 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.

Answer: B

Explanation:
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.
https://aws.amazon.com/about-aws/whats-new/2019/07/introducing-aws-budgets-reports/#:~:text=AWS%20Bud

NEW QUESTION 5
- (Exam Topic 1)
A company wants to host a new global website that consists of static content. A solutions architect is working on a solution that uses Amazon CloudFront with an origin access identity <OAI) to access website content that is stored in a private Amazon S3 bucket.
During testing, the solutions architect receives 404 errors from the S3 bucket. Error messages appear only for attempts to access paths that end with a forward slash. such as example.com/path/. These requests should return the existing S3 object path/index.html. Any potential solution must not prevent CloudFront from caching the content.
What should the solutions architect do to resolve this problem?

A. Change the CloudFront origin to an Amazon API Gateway proxy endpoin
B. Rewrite the S3 request URL by using an AWS Lambda function.
C. Change the CloudFront origin to an Amazon API Gateway endpoin
D. Rewrite the S3 request URL in an AWS service integration.
E. Change the CloudFront configuration to use an AWS Lambda@Edge function that is invoked by a viewer request event to rewrite the S3 request URL.
F. Change the CloudFront configuration to use an AWS Lambda@Edge function that is invoked by an origin request event to rewrite the S3 request URL.

Answer: C

NEW QUESTION 6
- (Exam Topic 1)
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 (or 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 markeling 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 SOS queue and trigger an AWS Lambda function lo process them.
C. Receive the orders using the AWS Step Functions program and trigger an Amazon ECS container lo process them.
D. Receive the orders in Amazon Kinesis Data Streams and use Amazon EC2 instances to process them.

Answer: B

Explanation:
Q: How does Amazon Kinesis Data Streams differ from Amazon SQS?
Amazon Kinesis Data Streams enables real-time processing of streaming big data. It provides ordering of records, as well as the ability to read and/or replay records in the same order to multiple Amazon Kinesis Applications. The Amazon Kinesis Client Library (KCL) delivers all records for a given partition key to the same record processor, making it easier to build multiple applications reading from the same Amazon Kinesis data stream (for example, to perform counting, aggregation, and filtering).
https://aws.amazon.com/kinesis/data-streams/faqs/
https://aws.amazon.com/blogs/big-data/unite-real-time-and-batch-analytics-using-the-big-data-lambda-architect

NEW QUESTION 7
- (Exam Topic 1)
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 onl
B. Use 3-year scheduled Reserved Instances for the web server EC2 instance
C. Detach the internet gateway and remove the NAT gateways from the VP
D. Use an Aurora Servertess database and set up a VPC endpoint for the S3 bucket.
E. Review the Auto Scaling group settings and ensure the scheduled actions are specified to operate the Amazon EC2 instances during business hours onl
F. Detach the internet gateway and remove the NAT gateways from the VP

G. Use an Aurora Servertess 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.
H. Review the Auto Scaling group settings and ensure the scheduled actions are specified to operate the Amazon EC2 instances during business hours onl
I. Detach the internet gateway from the VPC, and use an Aurora Servertess databas
J. Set up a VPC endpoint for the S3 bucket, then update the network routing and security rules and policies related to the changes.
K. Use 3-year scheduled Reserved Instances for the web server Amazon EC2 instance
L. Remove the NAT gateways from the VPC, and set up a VPC endpoint for the S3 bucke
M. Use Amazon
N. CloudWatch and AWS Lambda to stop and start the Aurora DB cluster so it operates during business hours onl
O. Update the network routing and security rules and policies related to the changes.

Answer: B

Explanation:
The application is accessible from the company network only remove NAT and IGW, application - S3 with VPC endpoint. Non-Production application no need to go for Reserved instances
To build site-to-site vpn, you don't need internet gateway. Instead, customer gateway is needed.
https://docs.aws.amazon.com/vpn/latest/s2svpn/SetUpVPNConnections.html#vpn-create-cgw

NEW QUESTION 8
- (Exam Topic 1)
A financial services company receives a regular data feed from its credit card servicing partner Approximately 5.1 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 queu
B. Trigger another Lambda function when new messages arrive in the SOS queue to process the records, writing the results to a temporary location in Amazon S3. Trigger a final Lambda function once the SOS queue is empty to transform the records into JSON format and send the results to another S3 bucket for internal processing.
C. Tigger an AWS Lambda function on file delivery that extracts each record and wntes it to an Amazon SOS queu
D. Configure an AWS Fargate container application to
E. automatically scale to a single instance when the SOS queue contains message
F. Have the application process each record, and transform the record into JSON forma
G. When the queue is empty, send the results to another S3 bucket for internal processing and scale down the AWS Fargate instance.
H. 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 requirement
I. Define the output format as JSO
J. Once complete, have the ETL job send the results to another S3 bucket for internal processing.
K. Create an AWS Glue crawler and custom classifier based upon the data feed formats and build a table definition to matc
L. 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 requirement
M. Define the output format as JSO
N. Once complete, send the results to another S3 bucket for internal processing and scale down the EMR cluster.

Answer: C

Explanation:
You can use a Glue crawler to populate the AWS Glue Data Catalog with tables. The Lambda function can be triggered using S3 event notifications when object create events occur. The Lambda function will then trigger the Glue ETL job to transform the records masking the sensitive data and modifying the output format to JSON. This solution meets all requirements.
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.
https://docs.aws.amazon.com/glue/latest/dg/trigger-job.html
https://d1.awsstatic.com/Products/product-name/diagrams/product-page-diagram_Glue_Event-driven-ETL-Pipel

NEW QUESTION 9
- (Exam Topic 1)
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 lo 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 connectio
C. Use inter-region VPC peering to access the 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 connectio
E. Use an AWS transit VPC solution to access data in other AWS Regions.
F. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region.Use the company WAN to send traffic over a DX connectio
G. Use Direct Connect Gateway to access data in other AWS Regions.

Answer: D

Explanation:
This feature also allows you to connect to any of the participating VPCs from any Direct Connect location, further reducing your costs for making using AWS services on a cross-region basis. https://aws.amazon.com/blogs/aws/new-aws-direct-connect-gateway-inter-region-vpc-access/
https://docs.aws.amazon.com/whitepapers/latest/aws-vpc-connectivity-options/aws-direct-connect-aws-transit-g

NEW QUESTION 10
- (Exam Topic 1)
A solutions architect is responsible (or 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 send within 10 seconds before they will time out and retry the request.
Which approach would improve the availability and durability of (he 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 functio
B. 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.
C. Create an Amazon API Gateway REST API that uses a service proxy to put items in an Amazon SOS queu
D. Extract the core processing code from the existing application and update it to pull items from Amazon SOS instead of an in-memory queu
E. 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 SOS queue.
F. Modify the application to use Amazon DynamoDB instead of Amazon RD
G. Configure Auto Scaling for the DynamoDB tabl
H. Deploy the application within an Auto Scaling group with a scaling policy based on CPU utilizatio
I. Back the in-memory queue with a memory-mapped file to an instance store volume and periodically write that file to Amazon S3.
J. Update the application to use a Redis task queue instead of the in-memory queu
K. 8uild a Docker container image for the applicatio
L. Create an Amazon ECS task definition that includes the application container and a separate container to host Redi
M. Deploy the new task definition as an ECS service using AWS Fargate, and enable Auto Scaling.

Answer: B

Explanation:
The obvious challenges here are long workloads, scalability based on queue load, and reliability. Almost always the defacto answer to queue related workload is SQS. Since the workloads are very long (90 minutes) Lambdas cannot be used (15 mins max timeout). So, autoscaled smaller EC2 nodes that wait on external services to complete the task makes more sense. If the task fails, the message is returned to the queue and retried.

NEW QUESTION 10
- (Exam Topic 1)
A company has developed a single-page web application in JavaScript. The source code is stored in a single Amazon S3 bucket in the us-east-1 Region. The company serves the web application to a global user base through Amazon CloudFront.
The company wants to experiment with two versions of the website without informing application users. Each version of the website will reside in its own S3 bucket. The company wants to determine which version is most successful in marketing a new product.
The solution must send application users that are based in Europe to the new website design. The solution must send application users that are based in the United States to the current website design. However, some exceptions exist. The company needs to be able to redirect specific users to the new website design, regardless of the users' location.
Which solution meets these requirements?

A. Configure two CloudFront distribution
B. Configure a geolocation routing policy in Amazon Route 53 to route traffic to the appropriate CloudFront endpoint based on the location of clients.
C. Configure a single CloudFront distributio
D. Create a behavior with different paths for each version of the sit
E. Configure Lambda@Edge on the default path to generate redirects and send the client to the correct version of the website.
F. Configure a single CloudFront distributio
G. Configure an alternate domain name on the distribution.Configure two behaviors to route users to the different S3 origins based on the domain name that the client uses in the HTTP request.
H. Configure a single CloudFront distribution with Lambda@Edg
I. Use Lambda@Edge to send user requests to different origins based on request attributes.

Answer: A

NEW QUESTION 11
- (Exam Topic 1)
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 solutions 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 Log
B. Use Amazon Athena to analyze the logs for traffic that can be remove
C. Ensure that security groups are blocking traffic that is responsible for high costs.
D. Add an interface VPC endpoint for Kinesis Data Streams to the VP
E. Ensure that applications have thecorrect IAM permissions to use the interface VPC endpoint.
F. Enable VPC Flow Logs and Amazon Detectiv
G. Review Detective findings for traffic that is not related to Kinesis Data Streams Configure security groups to block that traffic
H. Add an interface VPC endpoint for Kinesis Data Streams to the VPC Ensure that the VPC endpoint policy allows traffic from the applications

Answer: D

Explanation:
https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-endpoints-access.html https://aws.amazon.com/premiumsupport/knowledge-center/vpc-reduce-nat-gateway-transfer-costs/
VPC endpoint policies enable you to control access by either attaching a policy to a VPC endpoint or by using additional fields in a policy that is attached to an IAM user, group, or role to restrict access to only occur via the specified VPC endpoint

NEW QUESTION 15
- (Exam Topic 1)
A company manages an on-premises JavaScript front-end web application. The application is hosted on two servers secured with a corporate Active Directory. The application calls a set of Java-based microservices on an application server and stores data in a clustered MySQL database. The application is heavily used during the day on weekdays. It is lightly used during the evenings and weekends.
Daytime traffic to the application has increased rapidly, and reliability has diminished as a result. The company wants to migrate the application to AWS with a solution that eliminates the need for server maintenance, with an API to securely connect to the microservices.
Which combination of actions will meet these requirements? (Select THREE.)

A. Host the web application on Amazon S3. Use Amazon Cognito identity pools (federated identities) with SAML for authentication and authorization.
B. Host the web application on Amazon EC2 with Auto Scalin
C. Use Amazon Cognito federation and Login with Amazon for authentication and authorization.
D. Create an API layer with Amazon API Gatewa
E. Rehost the microservices on AWS Fargate containers.
F. Create an API layer with Amazon API Gatewa
G. Rehost the microservices on Amazon Elastic Container Service (Amazon ECS) containers.
H. Replatform the database to Amazon RDS for MySQL.
I. Replatform the database to Amazon Aurora MySQL Serverless.

Answer: ACE

NEW QUESTION 16
- (Exam Topic 1)
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 disable
B. Store statements in S3 Standar
C. Define an S3 Lifecycle policy to transition the data to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 day
D. Define another S3 Lifecycle policy to move the data to S3 Glacier Deep Archive after 2 year
E. Attach an S3 Glacier Vault Lock policy with deny delete permissions for archives less than 7 years old.
F. Create an S3 bucket with versioning enable
G. Store statements in S3 Intelligent-Tierin
H. Usesame-Region replication to replicate objects to a backup S3 bucke
I. Define an S3 Lifecycle policy for the backup S3 bucket to move the data to S3 Glacie
J. Attach an S3 Glacier Vault Lock policy with deny delete permissions for archives less than 7 years old.
K. Create an S3 bucket with Object Lock enable
L. Store statements in S3 Intelligent-Tierin
M. Enable compliance mode with a default retention period of 2 year
N. Define an S3 Lifecycle policy to move the data to S3 Glacier after 2 year
O. Attach an S3 Glacier Vault Lock policy with deny delete permissionsfor archives less than 7 years old.
P. Create an S3 bucket with versioning disable
Q. Store statements in S3 One Zone-Infrequent Access (S3 One Zone-IA). Define an S3 Lifecyde policy to move the data to S3 Glacier Deep Archive after 2 year
R. Attach an S3 Glader Vault Lock policy with deny delete permissions for archives less than 7 years old.

Answer: C

Explanation:
https://aws.amazon.com/about-aws/whats-new/2018/11/s3-object-lock/
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.
https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock-overview.html

NEW QUESTION 19
- (Exam Topic 1)
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:
* 1. The data must be highly durable and available.
* 2. The data must always be encrypted at rest and in transit.
* 3. 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 mod
B. Use Amazon EBS volume encryption using an AWS KMS key to encrypt the storage gateway volumes.
C. Use Amazon S3 with a bucket policy to enforce HTTPS for connections to the bucket and to enforce server-side encryption and AWS KMS for object encryption.
D. Use Amazon DynamoDB with SSL to connect to DynamoD
E. Use an AWS KMS key to encrypt DynamoDB objects at rest.
F. Deploy instances with Amazon EBS volumes attached to store this dat
G. Use E8S volume encryption using an AWS KMS key to encrypt the data.

Answer: B

Explanation:
Use Amazon S3 with a bucket policy to enforce HTTPS for connections to the bucket and to enforce server-side encryption and AWS KMS for object encryption.

NEW QUESTION 23
- (Exam Topic 1)
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 stor
B. Use VM Import/Export to move the VMs to Amazon EC2.
C. Use the VMware vSphere client to export the application as an image in Open Virealization Format (OVF) format Create an Amazon S3 bucket to store the image in the destination AWS Regio
D. Create and apply an IAM role for VM Import Use the AWS CLI to run the EC2 import command.
E. Configure AWS Storage Gateway for files service to export a Common Internet File System (CIFSJ shar
F. Create a backup copy to the shared folde
G. 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.
H. Create a managed-instance activation for a hybrid environment in AWS Systems Manage
I. 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 AM
J. Launch an EC2 instance that is based on the AMI

Answer: B

Explanation:
https://docs.aws.amazon.com/vm-import/latest/userguide/vmimport-image-import.html
- Export an OVF Template
- Create / use an Amazon S3 bucket for storing the exported images. The bucket must be in the Region where you want to import your VMs.
- Create an IAM role named vmimport.
- You'll use AWS CLI to run the import commands. https://aws.amazon.com/premiumsupport/knowledge-center/import-instances/

NEW QUESTION 26
- (Exam Topic 1)
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 Balance
B. Allocate an Amazon Workspaces Workspace for each end user to improve the user experience.
C. Migrate the database to an Amazon RDS Aurora PostgreSQL configuratio
D. Host the application and presentation layers in an Auto Scaling configuration on Amazon EC2 instances behind an Application Load Balance
E. Use Amazon AppStream 2.0 to improve the user experience.
F. Migrate the database to an Amazon RDS PostgreSQL Mulli-AZ configuratio
G. Host the application and presentation layers in automatically scaled AWS Fargate containers behind a Network Load Balance
H. Use Amazon ElastiCache to improve the user experience.
I. Migrate the database to an Amazon Redshift cluster with at least two node
J. Combine and host the application and presentation layers in automatically scaled Amazon ECS containers behind an Application Load Balance
K. Use Amazon CloudFront to improve the user experience.

Answer: B

Explanation:
Aurora would improve availability that can replicate to multiple AZ (6 copies). Auto scaling would improve the performance together with a ALB. AppStream is like Citrix that deliver hosted Apps to users.

NEW QUESTION 29
- (Exam Topic 1)
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 recovery 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 Regio
B. Use inter-Region VPC peering to connect both VPC
C. Deploy an Application Load Balancer (ALB) spanning multiple Availability Zones (AZs) to the VPC in theus-east-1 Regio
D. Deploy EC2 instances across multiple AZs in each Region as part of an Auto Scaling group spanning both VPCs and served by the ALB.
E. Deploy an Application Load Balancer (ALB) spanning multiple Availability Zones (AZs) to the VPC inthe us-east-1 Regio
F. Deploy EC2 instances across multiple AZs as part of an Auto Scaling group served by the AL
G. Deploy 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.
H. Create a VPC in the us-west-1 Regio
I. 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 AL
J. Create an Amazon Route 53 record that points to the ALB.

K. Deploy an Application Load Balancer (ALB) spanning multiple Availability Zones (AZs) to the VPC in the us-east-1 Regio
L. Deploy EC2 instances across multiple AZs as part of an Auto Scaling group served by the AL
M. Deploy the same solution to the us-west-1 Regio
N. Create separate Amazon Route 53 records in each Region that point to the ALB in the Regio
O. Use Route 53 health checks to provide high availability across both Regions.

Answer: B

Explanation:
A new web application in a active-passive DR mode. a Route 53 record set with a failover routing policy.

NEW QUESTION 30
- (Exam Topic 1)
A company wants to control its cost of Amazon Athena usage The company has allocated a specific monthly budget for Athena usage A solutions architect must design a solution that will prevent the company from exceeding the budgeted amount
Which solution will moot these requirements?

A. Use AWS Budget
B. Create an alarm (or when the cost of Athena usage reaches the budgeted amount for the mont
C. Configure AWS Budgets actions to deactivate Athena until the end of the month.
D. Use Cost Explorer to create an alert for when the cost of Athena usage reaches the budgeted amount for the mont
E. Configure Cost Explorer to publish notifications to an Amazon Simple Notification Service (Amazon SNS) topic.
F. Use AWS Trusted Advisor to track the cost of Athena usag
G. Configure an Amazon EventBridge (Amazon CloudWatch Events) rule to deactivate Athena until the end of the month whenever the cost reaches the budgeted amount for the month
H. Use Athena workgroups to set a limit on the amount of data that can be scanne
I. Set a limit that is appropriate for the monthly budget and the current pricing for Athena.

Answer: D

NEW QUESTION 35
- (Exam Topic 1)
A developer reports receiving an Error 403: Access Denied message when they try to download an object from an Amazon S3 bucket. The S3 bucket is accessed using an S3 endpoint inside a VPC. and is encrypted with an AWS KMS key. A solutions architect has verified that (he developer is assuming the correct IAM role in the account that allows the object to be downloaded. The S3 bucket policy and the NACL are also valid.
Which additional step should the solutions architect take to troubleshoot this issue?

A. Ensure that blocking all public access has not been enabled in the S3 bucket.
B. Verify that the IAM rote has permission to decrypt the referenced KMS key.
C. Verify that the IAM role has the correct trust relationship configured.
D. Check that local firewall rules are not preventing access to the S3 endpoint.

Answer: B

NEW QUESTION 37
- (Exam Topic 1)
A company uses AWS Transit Gateway for a hub-and-spoke model to manage network traffic between many VPCs. The company is developing a new service that must be able to send data at 100 Gbps. The company needs a faster connection to other VPCs in the same AWS Region.
Which solution will meet these requirements?

A. Establish VPC peering between the necessary VPC
B. Ensure that all route tables are updated as required.
C. Attach an additional transit gateway to the VPC
D. Update the route tables accordingly.
E. Create AWS Site-to-Site VPN connections that use equal-cost multi-path (ECMP) routing between the necessary VPCs.
F. Create an additional attachment from the necessary VPCs to the existing transit gateway.

Answer: D

NEW QUESTION 40
- (Exam Topic 1)
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 server
B. Use the AMIs to provision EC2 instances to recreate the applications in the AWS.Clou
C. Place an Application Load Balancer (ALB) in front of the EC2 instance
D. Use Amazon Route 53 to point the DNS names of the web forms to the ALB.
E. 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 item
F. Create an Amazon Kinesis data stream to receive the data input and store the input in DynamoD
G. Use Amazon Route 53 to point the DNS names of the web forms to the Kinesis data stream's endpoint.
H. Create Docker images for each server of the legacy web form application
I. Create an Amazon Elastic Container Service (Amazon ECS) cluster on AWS Fargat
J. Place an Application Load Balancer in front of the ECS cluste
K. Use Fargate task storage to store the web form data.
L. Provision an Amazon Aurora Serverless cluste

M. Build multiple schemas for each web form's data storag
N. Use Amazon API Gateway and an AWS Lambda function to recreate the data input form
O. Use Amazon Route 53 to point the DNS names of the web forms to their corresponding API Gateway endpoint.

Answer: D

Explanation:
Provision an Amazon Aurora Serverless cluster. Build multiple schemas for each web forms 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.

NEW QUESTION 44
- (Exam Topic 2)
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 RD
B. and AWS CloudTrail Load customer records in Amazon RDS MySQL and train users to run queries using the AWS CL
C. 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
D. 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 tram users to run queries using the CLI via Athena Analyze CloudTrail events to audit and alarm on queries against personal data
E. Apply a service control policy (SCP) that denies access to all services except IAM Amazon DynamoD
F. and AWS CloudTrail Store customer records in DynamoDB and train users to run 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
G. 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 run queries using the AWS CLI Enable S3 object-level logging and analyze CloudTrail events to audit and alarm on queries against personal data

Answer: B

NEW QUESTION 45
- (Exam Topic 2)
A company is running a two-tier web-based application in an on-premises data center. The application layer 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 Replica
B. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled
C. Enable Aurora Auto Scaling for Aurora writer
D. Use an Application Load Balancer with the round robin routing algorithm and sticky sessions enabled
E. Aurora Auto Scaling for Aurora Replica
F. Use an Application Load Balancer with the round robin routing algorithm and sticky sessions enabled.
G. Aurora Auto Scaling for Aurora writer
H. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.

Answer: C

NEW QUESTION 49
- (Exam Topic 2)
A company wants to migrate its workloads from on premises to AWS. The workloads run on Linux and Windows. The company has a large on-premises intra structure that consists of physical machines and VMs that host numerous applications.
The company must capture details about the system configuration. system performance. running processure and network coi.net lions of its o. -premises ,on boards. The company also must divide the on-premises applications into groups for AWS migrations. The company needs recommendations for Amazon EC2 instance types so that the company can run its workloads on AWS in the most cost-effective manner.
Which combination of steps should a solutions architect take to meet these requirements? (Select THREE.)

A. Assess the existing applications by installing AWS Application Discovery Agent on the physical machines and VMs.
B. Assess the existing applications by installing AWS Systems Manager Agent on the physical machines and VMs
C. Group servers into applications for migration by using AWS Systems Manager Application Manager.
D. Group servers into applications for migration by using AWS Migration Hub.
E. Generate recommended instance types and associated costs by using AWS Migration Hub.
F. Import data about server sizes into AWS Trusted Adviso
G. Follow the recommendations for cost optimization.

Answer: BDF

NEW QUESTION 50
- (Exam Topic 2)
A company is running a three-tier web application in an on-premises data center. The frontend is served by an Apache web server, the middle tier is a monolithic Java application, and the storage tier is a PostgreSOL database.
During a recent marketing promotion, customers could not place orders through the application because the application crashed An analysis showed that all three tiers were overloaded. The application became unresponsive, and the database reached its capacity limit because of read operations. The company already has several similar promotions scheduled in the near future.

A solutions architect must develop a plan for migration to AWS to resolve these issues. The solution must maximize scalability and must minimize operational effort.
Which combination of steps will meet these requirements? (Select THREE.)

A. Refactor the frontend so that static assets can be hosted on Amazon S3. Use Amazon CloudFront to serve the frontend to customer
B. Connect the frontend to the Java application.
C. Rehost the Apache web server of the frontend on Amazon EC2 instances that are in an Auto Scaling grou
D. Use a load balancer in front of the Auto Scaling grou
E. Use Amazon Elastic File System (Amazon EFS) to host the static assets that the Apache web server needs.
F. Rehost the Java application in an AWS Elastic Beanstalk environment that includes auto scaling.
G. Refactor the Java applicatio
H. Develop a Docker container to run the Java applicatio
I. Use AWS Fargate to host the container.
J. Use AWS Database Migration Service (AWS DMS) to replatform the PostgreSQL database to an Amazon Aurora PostgreSQL databas
K. Use Aurora Auto Scaling for read replicas.
L. Rehost the PostgreSQL database on an Amazon EC2 instance that has twice as much memory as the on-premises server.

Answer: BCF

NEW QUESTION 53
- (Exam Topic 2)
A company is migrating an on-premises application and a MySQL database to AWS. The application processes highly sensitive data, and new data is constantly updated in the database. The data must not be transferred over the internet. The company also must encrypt the data in transit and at rest.
The database is 5 TB in size. The company already has created the database schema in an Amazon RDS for MySQL DB instance The company has set up a 1 Gbps AWS Direct Connect connection to AWS. The company also has set up a public VIF and a private VIF. A solutions architect needs to design a solution that will migrate the data to AWS with the least possible downtime
Which solution will meet these requirements?

A. Perform a database backu
B. Copy the backup files to an AWS Snowball Edge Storage Optimized device.Import the backup to Amazon S3. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3) for encryption at rest Use TLS for encryption in transit Import the data from Amazon S3 to the DB instance.
C. Use AWS Database Migration Service (AWS DMS) to migrate the data to AW
D. Create a DMS replication instance in a private subne
E. Create VPC endpoints for AWS DM
F. Configure a DMS task to copy data from the on-premises database to the DB instance by using full load plus change data capture (CDC). Use the AWS Key Management Service (AWS KMS) default key for encryption at res
G. Use TLS for encryption in transit.
H. Perform a database backu
I. Use AWS DataSync to transfer the backup files to Amazon S3 Useserver-side encryption with Amazon S3 managed encryption keys (SSE-S3) for encryption at res
J. Use TLS for encryption in transit Import the data from Amazon S3 to the DB instance.
K. Use Amazon S3 File Gateway Set up a private connection to Amazon S3 by using AWS PrivateLink.Perform a database backu
L. Copy the backup files to Amazon S3. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3) for encryption at res
M. Use TLS for encryption in transi
N. Import the data from Amazon S3 to the DB instance.

Answer: D

NEW QUESTION 55
- (Exam Topic 2)
A company wants to use Amazon Workspaces in combination with thin client devices to replace aging desktops Employees use the desktops to access applications that work with clinical trial data Corporate security policy states that access to the applications must be restricted to only company branch office locations. The company is considering adding an additional branch office in the next 6 months.
Which solution meets these requirements with the MOST operational efficiency?

A. Create an IP access control group rule with the list of public addresses from the branch offices Associate the IP access control group with the Workspaces directory
B. Use AWS Firewall Manager to create a web ACL rule with an IPSet with the list of public addresses from the branch office locations Associate the web ACL with the Workspaces directory
C. Use AWS Certificate Manager (ACM) to issue trusted device certificates to the machines deployed in the branch office locations Enable restricted access on the Workspaces directory
D. Create a custom Workspace image with Windows Firewall configured to restrict access to the public addresses of the branch offices Use the image to deploy the Workspaces.

Answer: C

NEW QUESTION 56
- (Exam Topic 2)
A company is planning to migrate an application from on premises to the AWS Cloud. The company will begin the migration by moving the application's underlying data storage to AWS The application data is stored on a shared tie system on premises, and the application servers connect to the shared We system through SMB.
A solutions architect must implement a solution that uses an Amazon S3 bucket tor shared storage Until the application Is fully migrated and code is rewritten to use native Amazon S3 APIs, the application must continue to have access to the data through SMB The solutions architect must migrate the application data to AWS to its new location while still allowing the on-premises application to access the data.
Which solution will meet these requirements?

A. Create a new Amazon FSx for Windows File Server fie system Configure AWS DataSync with one location tor the on-premises file share and one location for the new Amazon FSx file system Create a new DataSync task to copy the data from the on-premises file share location to the Amazon FSx file system
B. Create an S3 bucket for the applicatio
C. Copy the data from the on-premises storage to the S3 bucket

D. Deploy an AWS Server Migration Service (AWS SMS) VM to the on-premises environmen
E. Use AWS SMS to migrate the file storage server from on premises to an Amazon EC2 instance
F. Create an S3 bucket for the applicatio
G. Deploy a new AWS Storage Gateway Me gateway on anon-premises V
H. Create a new file share that stores data in the S3 bucket and is associated with the tie gatewa
I. Copy the data from the on-premises storage to the new file gateway endpoint.

Answer: A

NEW QUESTION 61
- (Exam Topic 2)
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 stat c 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 Auto Scaling group that processes the SQS queu
B. Replace the custom software with Amazon Recognition to categorize the videos.
C. Store the uploaded videos n Amazon EFS and mount the file system to the EC2 instances for Te web applicatio
D. Process the SOS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos.
E. Host the web application in Amazon S3. Store the uploaded videos in Amazon S3. Use S3 event notifications to publish events to the SQS queue Process the SQS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos.
F. Use AWS Elastic Beanstalk to launch EC2 instances in an Auto Scaling group for the web application and launch a worker environment to process the SQS queue Replace the custom software with Amazon Rekognition to categorize the videos.

Answer: D

NEW QUESTION 62
- (Exam Topic 2)
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 tor 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 m the identity accoun
B. Create roles with appropriate permissions in the production and testing accounts Add the identity account to the trust policies for the roles
C. 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.
D. 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.
E. Create all user accounts in the production account Create roles for access in me production account and testing account
F. Grant cross-account access from the production account to the testing account

Answer: A

NEW QUESTION 63
- (Exam Topic 2)
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 toeu-west-1.

Answer: B

NEW QUESTION 64
- (Exam Topic 2)
A solutions architect uses AWS Organizations to manage several AWS accounts for a company. The full Organizations feature set is activated for the organization. All production AWS accounts exist under an OU that is named "production ‘’ Systems operators have full administrative privileges within these accounts by using IAM roles.
The company wants to ensure that security groups in all production accounts do not allow inbound traffic for TCP port 22. All noncompliant security groups must be remediated immediately, and no new rules that allow port 22 can be created.
Winch solution will meet these requirements?

A. Write an SCP that denies the CreateSecurityGroup action with a condition o(ec2:tngress rule with value 22. Apply the SCP to the 'production' OU.
B. Configure an AWS CloudTrail trail for all accounts Send CloudTrail logs to an Amazon S3 bucket In the Organizations management accoun
C. Configure an AWS Lambda function on the management account with permissions to assume a role in all production accounts to describe and modify security group
D. Configure Amazon S3 to invoke the Lambda function on every PutObject event on the S3 bucket Configure the Lambda function to analyze each CloudTrail

event for noncompliant security group actions and to automatically remediate any issues.
E. Create an Amazon EvertBridge (Amazon CloudWatch Events) event bus in the Organizations management accoun
F. Create an AWS Cloud Formation template to deploy configurations that send CreateSecurityGroup events to the even! bus from an production accounts Configure an AWS Lambda function in the management account with permissions to assume a role «i all production accounts to describe and modify security group
G. Configure the event bus to invoke the Lambda function Configure the Lambda function to analyse each event for noncompliant security group actions and to automatically remediate any issues.
H. Create an AWS CloudFormation template to turn on AWS Config Activate the INCOMING_SSH_DISABLED AWS Config managed rule Deploy an AWS Lambda function that will run based on AWS Config findings and will remediate noncompliant resources Deploy the CloudFormation template by using a StackSet that is assigned to the "production" O
I. Apply an SCP to the OU to deny modification of the resources that the CloudFormation template provisions.

Answer: D

NEW QUESTION 65
- (Exam Topic 2)
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 connectio
B. Create a new SMB file shar
C. Write nightly database exports to the new SMB file share.
D. Create an Amzon FSx for Windows File Server Single-AZ file system within the VPC that is connected to the Direct Connect connectio
E. Create a new SMB file shar
F. Write nightly database exports to an SMB file share on the Amazon FSx file system Enable backups.
G. Create an Amazon FSx for Windows File Server Multi-AZ system within the VPC that is connected to the Direct Connect connectio
H. Create a new SMB file shar
I. Write nightly database exports to an SMB file share on the Amazon FSx file syste
J. Enable nightly backups.
K. Create a new S3 bucket
L. Deploy an AWS Storage Gateway volume gateway within the VPC that is connected to the Direct Connect connectio
M. Create a new SMB file shar
N. Write nightly database exports to the new SMB file share on the volume gateway, and automate copies of this data to an S3 bucket.

Answer: A

NEW QUESTION 66
- (Exam Topic 2)
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 Cloud Formation 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 rote.
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

Answer: B

NEW QUESTION 69
- (Exam Topic 2)
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 an AWS accounts. The company plans to use Amazon Elasticsearch Service (Amazon ES) to perform log analyses in me logging account.
Which strategy should a solutions architect use to meet These requirements?

A. Configure CloudTrail and VPC Flow Logs m each AWS account to send data to a centralized Amazon S3 Ducket in the fogging accoun
B. Create an AWS Lambda function to load data from the S3 bucket to Amazon ES m the togging account
C. Configure CloudTrail and VPC Flow Logs to send data to a fog group m Amazon CloudWatch Logs n each AWS account Configure a CloudWatch subscription filter m each AWS account to send data to Amazon Kinesis Data Firehose In the fogging account Load data from Kinesis Data Firehose Into Amazon ES in the logging account
D. Configure CloudTrail and VPC Flow Logs to send data to a separate Amazon S3 bucket In each AWS accoun
E. Create an AWS Lambda function triggered by S3 evens to copy the data to a centralized logging bucke
F. Create another Lambda function lo load data from the S3 bucket to Amazon ES in the logging account.
G. Configure CloudTrail and VPC Flow Logs to send data to a fog group in Amazon CloudWatch Logs n each AWS account Create AWS Lambda functions in each AWS account to subscribe to the tog groups and stream the data to an Amazon S3 bucket in the togging accoun
H. Create another Lambda function to toad data from the S3 bucket to Amazon ES in the logging account.

Answer: A

NEW QUESTION 70

- (Exam Topic 2)
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 a low developers the possibilityy 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 StrmgEquals condition for the DevelopmentUnit resource tag and aws Principal Tag '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

Answer: A

NEW QUESTION 74
- (Exam Topic 2)
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 Account.
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? (Select THREE.)

A. Create a bucket policy that includes read permissions for the S3 bucke
B. Set the principal of the bucket policy to the account ID of the Strategy account
C. Update the strategy_reviewer IAM role to grant full permissions for the S3 bucket and to grant decrypt permissions for the custom KMS key.
D. Update the custom KMS key policy in the Creative account to grant decrypt permissions to the strategy_reviewer IAM role.
E. Create a bucket policy that includes read permissions for the S3 bucke
F. Set the principal of the bucket policy to an anonymous user.
G. Update the custom KMS key policy in the Creative account to grant encrypt permissions to the strategy_reviewer IAM role.
H. Update the strategy_reviewer IAM role to grant read permissions for the S3 bucket and to grant decrypt permissions for the custom KMS key

Answer: ACE

NEW QUESTION 77
- (Exam Topic 2)
A company is running a critical application that uses an Amazon RDS for MySQL database to store data. The RDS DB instance is deployed in Multi-AZ mode.
A recent RDS database failover test caused a 40-second outage to the application A solutions architect needs to design a solution to reduce the outage time to less than 20 seconds.
Which combination of steps should the solutions architect take to meet these requirements? (Select THREE.)

A. Use Amazon ElastiCache for Memcached in front of the database
B. Use Amazon ElastiCache for Redis in front of the database.
C. Use RDS Proxy in front of the database
D. Migrate the database to Amazon Aurora MySQL
E. Create an Amazon Aurora Replica
F. Create an RDS for MySQL read replica

Answer: ABF

NEW QUESTION 78
- (Exam Topic 2)
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 Region solution that wifi 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 S3Cross-Region Replication Create a new DynamoDB able in a new Region Use the new table as a replica target tor DynamoDB global tables.
B. Create an Amazon CloudFront distribution to serve assets from the S3 bucke
C. Configure S3Same-Region Replicatio
D. Create a new DynamoDB able m a new Regio
E. Configure asynchronous replication between the DynamoDB tables by using AWS Database Migration Service (AWS DMS) with change data capture (CDC)
F. 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 Regio
G. Configure DynamoDB global tables by enabling Amazon DynamoDB Streams, and add a replica table in a new Region.
H. Create another S3 bucket in the same Region, and configure S3 Same-Region Replication between the buckets- Create an Amazon CloudFront distribution and

configure origin failover with two origin accessing the S3 buckets Create a new DynamoDB table m a new Region Use the new table as a replica target for DynamoDB global tables.

Answer: B

NEW QUESTION 82
- (Exam Topic 2)
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

Answer: A

NEW QUESTION 83
- (Exam Topic 2)
A company is configuring connectivity to a multi-account AWS environment to support application workloads fiat serve users in a single geographic region. The workloads depend on a highly available, on-premises legacy system deployed across two locations It is critical for the AWS workloads to manias 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 (OX) 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 me 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 VP
C. Create a DX gateway m 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 me DX gateway.
D. 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 accoun
E. Create a transit virtual interface for each DX interlace and associate them with the DX gatewa
F. Create a gateway association between the DX gateway and the transit gateway
G. 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 VP
H. 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.

Answer: B

NEW QUESTION 87
- (Exam Topic 2)
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 s3CreateAccessPoint 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: Create Access Point action unless the s3:AccessPointNetworkOrigin condition key evaluates to VPC.
C. Use AWS Cloud Formation StackSets to create a new IAM policy In each AWS account that allows the s3:CreateAccessPoint action only if the s3:AccessPointNetwofkOngm condition key evaluates to VPC.
D. Set the S3 bucket policy to deny the s3:CreateAccessPoint action unless the s3: AccessPointNetworkOngin condition key evaluates to VPC.

Answer: A

NEW QUESTION 88
- (Exam Topic 2)
A company runs many workloads on AWS and uses AWS Organizations to manage its accounts. The workloads are hosted on Amazon EC2, AWS Fargate, and AWS Lambda. Some of the workloads have unpredictable demand. Accounts record high usage in some months and low usage in other months.
The company wants to optimize its compute costs over the next 3 years. A solutions architect obtains a 6-month average for each of the accounts across the organization to calculate usage.
Which solution will provide the MOST cost savings for all the organization's compute usage?

A. Purchase Reserved Instances for the organization to match the size and number of the most common EC2 instances from the member accounts.
B. Purchase a Compute Savings Plan for the organization from the management account by using the recommendation at the management account level.
C. Purchase Reserved Instances for each member account that had high EC2 usage according to the data from the last 6 months.
D. Purchase an EC2 Instance Savings Plan for each member account from the management account based on EC2 usage data from the last 6 months.

Answer: A

NEW QUESTION 92
- (Exam Topic 2)

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 poin
C. Mount the file gateway file system on the on-premises serve
D. Configure a process to periodically copy the images to the mount point
E. 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 a 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
F. 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.

Answer: A

NEW QUESTION 96
- (Exam Topic 2)
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? (Select THREE)

A. Deploy two firewall appliances into the shared services VP
B. each in a separate Availability Zone
C. 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.
D. Create a new Gateway Load Balancer in the shared services VPC Create a new target group, and attach it to the new Gateway Load Balancer Add each of the firewall appliance instances to the target group
E. Create a VPC interface endpoint Add a route to the route table in the shared services VP
F. Designate the new endpoint as the next hop for traffic that enters the shared services VPC from other VPCs.
G. Deploy two firewall appliances into the shared services VP
H. each in the same Availability Zone

Answer: AC

NEW QUESTION 97
- (Exam Topic 2)
A solutions architect must update an application environment within AWS Elastic Beanstalk using a With 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

Answer: B

NEW QUESTION 101
- (Exam Topic 2)
A company has a platform that contains an Amazon S3 bucket for user content. The S3 bucket has thousands of terabytes of objects, all in the S3 Standard storage class. The company has an RTO of 6 hours The company must replicate the data from its primary AWS Region to a replication S3 bucket in another Region
The user content S3 bucket contains user-uploaded files such as videos and photos. The user content S3 bucket has an unpredictable access pattern. The number of users is increasing quickly, and the company wants to create an S3 Lifecycle policy to reduce storage costs
Which combination of steps will meet these requirements MOST cost-effectively'? (Select TWO)

A. Move the objects in the user content S3 bucket to S3 Intelligent-Tiering immediately
B. Move the objects in the user content S3 bucket to S3 Intelligent-Tiering after 30 days
C. Move the objects in the replication S3 bucket to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days and to S3 Glacier after 90 days
D. Move the objects in the replication S3 bucket to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 30 days and to S3 Glacier Deep Archive after 90 days
E. Move the objects in the replication S3 bucket to S3 Standard-infrequent Access (S3 Standard-IA) after 30 days and to S3 Glacier Deep Archive after 180 days

Answer: AD

NEW QUESTION 102
- (Exam Topic 2)
A company runs its application in the eu-west-1 Region and has one account for each of its environments development, testing, and production All the environments are running 24 hours a day 7 days a week by using stateful Amazon EC2 instances and Amazon RDS for MySQL databases The databases are between 500 GB and 800 GB in size
The development team and testing team work on business days during business hours, but the production environment operates 24 hours a day. 7 days a week. The company wants to reduce costs AH resources are tagged with an environment tag with either development, testing, or production as the key.

What should a solutions architect do to reduce costs with the LEAST operational effort?

A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs once every day Configure the rule to invoke one AWS Lambda function that starts or stops instances based on the tag day and time.
B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs every business day in the evenin
C. Configure the rule to invoke an AWS Lambda function that stops instances based on thetag-Create a second EventBridge (CloudWatch Events) rule that runs every business day in the morning Configure the second rule to invoke another Lambda function that starts instances based on the tag
D. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs every business day in the evening Configure the rule to invoke an AWS Lambda function that terminates instances based on the tag Create a second EventBridge (CloudWatch Events) rule that runs every business day in the morning Configure the second rule to invoke another Lambda function that restores the instances from their last backup based on the tag.
E. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs every hour Configure the rule to invoke one AWS Lambda function that terminates or restores instances from their ....based on the ta
F. day, and time

Answer: C

NEW QUESTION 103
- (Exam Topic 2)
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? (Select 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

Answer: ACF

NEW QUESTION 104
- (Exam Topic 2)
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-premise
B. 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
C. 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
D. 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
E. Set up an Amazon EC2 instance that runs the OpenCV libranes Copy the videos, images, and facecatalog from the on-premises library into an Amazon EBS volumemounted 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

Answer: C

NEW QUESTION 107
- (Exam Topic 2)
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.

Answer: B

NEW QUESTION 108
- (Exam Topic 2)
A solutions architect needs to review the design of an Amazon EMR cluster that is using the EMR File System (EMRFS). The cluster performs tasks that are critical to business needs. The cluster is running Amazon EC2 On-Demand Instances at all times for all task, master, and core nodes The EMR tasks run each

morning, starting at 1:00 AM, and take 6 hours to finish running. The amount of time to complete the processing is not a priority because the data is not referenced until late in the day.
The solutions architect must review the architecture and suggest a solution to minimize the compute costs Which solution should the solutions architect recommend to meet these requirements?

A. Launch all task, master, and core nodes on Spot Instances in an instance flee
B. Terminate the cluster, including all instances, when the processing is completed.
C. Launch the master and core nodes on On-Demand Instance
D. Launch the task nodes on Spot Instances In an instance flee
E. Terminate the cluster, including all instances, when the processing is complete
F. Purchase Compute Savings Plans to cover the On-Demand Instance usage.
G. Continue to launch all nodes on On-Demand Instance
H. Terminate the cluste
I. Including all instances, when the processing Is complete
J. Purchase Compute Savings Plans to cover the On-Demand Instance usage.
K. Launch the master and core nodes on On-Demand Instance
L. Launch the task nodes on Spot Instances In an instance flee
M. Terminate only the task node Instances when the processing is completed Purchase Compute Savings Plans to cover the On-Demand Instance usage.

Answer: B

NEW QUESTION 113
- (Exam Topic 2)
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 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? (Select TWO.)

A. Deploy the SOS 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 topics in each Region.
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.

Answer: CD

NEW QUESTION 118
- (Exam Topic 2)
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 microservic
B. Configure the API methods to require the key.
C. 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.
D. Modify the API Gateway to use IAM authentication Update the IAM policy for the IAM role that isassigned to the EC2 instances to allow access to the API Gateway Move the API Gateway into a new VPC Deploy a transit gateway and connect the VPCs.
E. 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 addres
F. Add an API key for each service to use for authentication.

Answer: B

NEW QUESTION 120
- (Exam Topic 2)
A company wants to deploy an API to AWS. The company plans to run the API on AWS Fargate behind a load balancer. The API requires the use of header-based routing and must be accessible from on-premises networks through an AWS Direct Connect connection and a private VIF.
The company needs to add the client IP addresses that connect to the API to an allow list in AWS. The company also needs to add the IP addresses of the API to the allow list. The company's security team will allow /27 CIDR ranges to be added to the allow list. The solution must minimize complexity and operational overhead.
Which solution will meet these requirements?

A. Create a new Network Load Balancer (NLB) in the same subnets as the Fargate task deployments.Create a security group that includes only the client IP addresses that need access to the AP
B. Attach the new security group to the Fargate task
C. Provide the security team with the NLB's IP addresses for the allow list.
D. Create two new /27 subnet
E. Create a new Application Load Balancer (ALB) that extends across the new subnet
F. Create a security group that includes only the client IP addresses that need access to the AP
G. Attach the security group to the AL
H. Provide the security team with the new subnet IP ranges for the allow list.
I. Create two new '27 subnet
J. Create a new Network Load Balancer (NLB) that extends across the new subnet
K. Create a new Application Load Balancer (ALB) within the new subnet
L. Create a security group that includes only the client IP addresses that need access to the AP

M. Attach the security group to the AL
N. Add the ALB's IP addresses as targets behind the NL
O. Provide the security team with the NLB's IP addresses for the allow list.
P. Create a new Application Load Balancer (ALB) in the same subnets as the Fargate task deployments.Create a security group that includes only the client IP addresses that need access to the AP
Q. Attach the security group to the AL
R. Provide the security team with the ALB's IP addresses for the allow list.

Answer: A

NEW QUESTION 122
- (Exam Topic 2)
A data analytics company has an Amazon Redshift cluster that consists of several reserved nodes. The duster 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 at) 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 duster 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 duster's CPU metrics in Amazon CloudWatch reach 80%.
C. Deploy an AWS Lambda function to add capacity to the Amazon Redshift duster by using an elastic resize operation when the duster's CPU metrics in Amazon CloudWatch leach 80%.
D. Turn on the Concurrency Scaling feature for the Amazon Redshift duster

Answer: D

NEW QUESTION 124
- (Exam Topic 2)
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.

Answer: A

NEW QUESTION 125
- (Exam Topic 2)
A company uses AWS Cloud Formation to deploy applications within multiple VPCs that are all attached to a transit gateway. Each VPC that sends traffic to the public internet must send the traffic through a shared services VPC. Each subnet within a VPC uses the default VPC route table, and the traffic is routed to the transit gateway. The transit gateway uses its default route table for any VPC attachment.
A security audit reveals that an Amazon EC2 instance that is deployed within a VPC can communicate with an EC2 instance that is deployed in any of the company's other VPCs. A solutions architect needs to limit the traffic between the VPCs. Each VPC must be able to communicate only with a predefined, limited set of authorized VPCs.
What should the solutions architect do to meet these requirements?

A. Update the network ACL of each subnet within a VPC to allow outbound traffic only to the authorized VPC
B. Remove all deny rules except the default deny rule.
C. Update all the security groups that are used within a VPC to deny outbound traffic to security groups that are used within the unauthorized VPCs
D. Create a dedicated transit gateway route table for each VPC attachmen
E. Route traffic only to the authorized VPCs.
F. Update the main route table of each VPC to route traffic only to the authorized VPCs through the transit gateway.

Answer: A

NEW QUESTION 127
- (Exam Topic 2)
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? (Select 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

Answer: CD

NEW QUESTION 128
- (Exam Topic 2)
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 endpomts as AWS Lambda functions behind a Network Load Balancer (NLB). and use Amazon RDS Proxy to wnte to the Aurora MySQL database When the databases are synced disable the replication job and restart the Aurora Replica as the primary instanc
B. Point the collector DNS record to the NLB.
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 Move the aggregation jobs to run against the Aurora MySQL database Set up collection endpomts 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 ALB Disable the AWS DMS sync task after the cutover from on premises to AWS
D. 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
E. 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.

Answer: C

NEW QUESTION 130
- (Exam Topic 2)
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:

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

A. Kms:GenerateDataKey
B. KmsGetKeyPolpcy
C. kmsGetPubKKey
D. kms:SKjn

Answer: A

NEW QUESTION 132
- (Exam Topic 2)
A company is migrating an on-premises content management system (CMS) to AWS Fargate. The company uses the CMS for blog posts that include text, images, and videos. The company has observed that traffic to blog posts drops by more than 80% after the posts are more than 30 days old
The CMS runs on multiple VMs and stores application state on disk This application state is shared across all instances across multiple Availability Zones Images and other media are stored on a separate NFS file share. The company needs to reduce the costs of the existing solution while minimizing the impact on performance.
Which combination of steps will meet these requirements MOST cost-effectively? (Select TWO.)

A. Store media in an Amazon S3 Standard bucket Create an S3 Lifecycle configuration that transitions objects that are older than 30 days to the S3 Standard-Infrequent Access (S3 Standard-IA) storage class.
B. Store media on an Amazon Elastic File System (Amazon EFS) volume Attach the EFS volume to all Fargate instances.
C. Store application state on an Amazon Elastic File System (Amazon EFS) volume Attach the EFS volume to all Fargate instances.
D. Store application state on an Amazon Elastic Block Store (Amazon EBS) volume Attach the EBS volume to all Fargate instances.
E. Store media in an Amazon S3 Standard bucket Create an S3 Lifecycle configuration that transitions objects that are older than 30 days to the S3 Glacier storage class

Answer: AC

NEW QUESTION 133
- (Exam Topic 2)
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

Answer: C

NEW QUESTION 138
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

SAP-C02 Practice Exam Features:

* SAP-C02 Questions and Answers Updated Frequently

* SAP-C02 Practice Questions Verified by Expert Senior Certified Staff

* SAP-C02 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* SAP-C02 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year