Exam Questions SAP-C02

AWS Certified Solutions Architect - Professional

https://www.2passeasy.com/dumps/SAP-C02/

NEW QUESTION 1
- (Exam Topic 1)
A company is running an application in the AWS Cloud. The company's security team must approve the creation of all new IAM users. When a new IAM user is created, all access for the user must be removed automatically. The security team must then receive a notification to approve the user. The company has a multi-Region AWS CloudTrail trail In the AWS account.
Which combination of steps will meet these requirements? (Select THREE.)

A. Create an Amazon EventBridge (Amazon CloudWatch Events) rul
B. Define a pattern with the detail-type value set to AWS API Call via CloudTrail and an eventName of CreateUser.
C. Configure CloudTrail to send a notification for the CreateUser event to an Amazon Simple Notification Service (Amazon SNS) topic.
D. Invoke a container that runs in Amazon Elastic Container Service (Amazon ECS) with AWS Fargate technology to remove access
E. Invoke an AWS Step Functions state machine to remove access.
F. Use Amazon Simple Notification Service (Amazon SNS) to notify the security team.
G. Use Amazon Pinpoint to notify the security team.

Answer: ADE

Explanation:
https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/send-a-notification-when-an-iam-user-is-crea

NEW QUESTION 2
- (Exam Topic 1)
A company has 50 AWS accounts that are members of an organization in AWS Organizations Each account contains multiple VPCs The company wants to use AWS Transit Gateway to establish connectivity between the VPCs in each member account Each time a new member account is created, the company wants to automate the process of creating a new VPC and a transit gateway attachment.
Which combination of steps will meet these requirements? (Select TWO)

A. From the management account, share the transit gateway with member accounts by using AWS Resource Access Manager
B. Prom the management account, share the transit gateway with member accounts by using an AWS Organizations SCP
C. Launch an AWS CloudFormation stack set from the management account that automatical^/ creates a new VPC and a VPC transit gateway attachment in a member accoun
D. Associate the attachment with the transit gateway in the management account by using the transit gateway ID.
E. Launch an AWS CloudFormation stack set from the management account that automatical^ creates a new VPC and a peering transit gateway attachment in a member accoun
F. Share the attachment with the transit gateway in the management account by using a transit gateway service-linked role.
G. From the management account, share the transit gateway with member accounts by using AWS Service Catalog

Answer: AC

Explanation:
https://aws.amazon.com/blogs/mt/self-service-vpcs-in-aws-control-tower-using-aws-service-catalog/ https://docs.aws.amazon.com/vpc/latest/tgw/tgw-transit-gateways.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-ec2-transitgatewayattachme

NEW QUESTION 3
- (Exam Topic 1)
A company manages multiple AWS accounts by using AWS Organizations. Under the root OU. the company has two OUs: Research and DataOps.
Because of regulatory requirements, all resources that the company deploys in the organization must reside in the ap-northeast-1 Region. Additionally. EC2 instances that the company deploys in the DataOps OU must use a predefined list of instance types
A solutions architect must implement a solution that applies these restrictions. The solution must maximize operational efficiency and must minimize ongoing maintenance
Which combination of steps will meet these requirements? (Select TWO)

A. Create an IAM role in one account under the DataOps OU Use the ec2 Instance Type condition key in an inline policy on the role to restrict access to specific instance types.
B. Create an IAM user in all accounts under the root OU Use the aws RequestedRegion condition key in an inline policy on each user to restrict access to all AWS Regions except ap-northeast-1.
C. Create an SCP Use the aws:RequestedRegion condition key to restrict access to all AWS Regions except ap-northeast-1 Apply the SCP to the root OU.
D. Create an SCP Use the ec2Reo»on condition key to restrict access to all AWS Regions except ap-northeast-1. Apply the SCP to the root O
E. the DataOps O
F. and the Research OU.
G. Create an SCP Use the ec2:InstanceType condition key to restrict access to specific instance types Apply the SCP to the DataOps OU.

Answer: CE

Explanation:
https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_examples_aws_deny-requested-region.h
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_ec2.html

NEW QUESTION 4
- (Exam Topic 1)
A company has an application that runs on Amazon EC2 instances. A solutions architect is designing VPC infrastructure in an AWS Region where the application needs to access an Amazon Aurora DB cluster. The EC2 instances are all associated with the same security group. The DB cluster is associated with its own security group.
The solutions architect needs to add rules to the security groups to provide the application with least privilege access to the DB cluster.
Which combination of steps will meet these requirements? (Select TWO.)

A. Add an inbound rule to the EC2 instances' security grou

B. Specify the DB cluster's security group as the source over the default Aurora port.
C. Add an outbound rule to the EC2 instances' security grou
D. Specify the DB cluster's security group as the destination over the default Aurora port.
E. Add an inbound rule to the DB cluster's security grou
F. Specify the EC2 instances' security group as the source over the default Aurora port.
G. Add an outbound rule to the DB cluster's security grou
H. Specify the EC2 instances' security group as the destination over the default Aurora port.
I. Add an outbound rule to the DB cluster's security grou
J. Specify the EC2 instances' security group as the destination over the ephemeral ports.

Answer: AB

Explanation:
* B. Add an outbound rule to the EC2 instances' security group. Specify the DB cluster's security group as the destination over the default Aurora port. This allows the instances to make outbound connections to the DB cluster on the default Aurora port. C. Add an inbound rule to the DB cluster's security group. Specify the EC2 instances' security group as the source over the default Aurora port. This allows connections to the DB cluster from the EC2 instances on the default Aurora port.

NEW QUESTION 5
- (Exam Topic 1)
An enterprise company wants to allow its developers to purchase third-party software through AWS Marketplace. The company uses an AWS Organizations account structure with full features enabled, and has a shared services account in each organizational unit (OU) that will be used by procurement managers. The procurement team's policy indicates that developers should be able to obtain third-party software from an approved list only and use Private Marketplace in AWS Marketplace to achieve this requirement . The procurement team wants administration of Private Marketplace to be restricted to a role named procurement-manager-role, which could be assumed by procurement managers Other IAM users groups, roles, and account administrators in the company should be denied Private Marketplace administrative access
What is the MOST efficient way to design an architecture to meet these requirements?

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

Answer: C

Explanation:
SCP to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role.
https://aws.amazon.com/blogs/awsmarketplace/controlling-access-to-a-well-architected-private-marketplace-usi
This approach allows the procurement managers to assume the procurement-manager-role in shared services accounts, which have the AWSPrivateMarketplaceAdminFullAccess managed policy attached to it and can then manage the Private Marketplace. The organization root-level SCP denies the permission to administer Private Marketplace to everyone except the role named procurement-manager-role and another SCP denies the permission to create an IAM role named procurement-manager-role to everyone in the organization, ensuring that only the procurement team can assume the role and manage the Private Marketplace. This approach provides a centralized way to manage and restrict access to Private Marketplace while maintaining a high level of security.

NEW QUESTION 6
- (Exam Topic 1)
A company has developed APIs that use Amazon API Gateway with Regional endpoints. The APIs call AWS Lambda functions that use API Gateway authentication mechanisms. After a design review, a solutions architect identifies a set of APIs that do not require public access.
The solutions architect must design a solution to make the set of APIs accessible only from a VPC. All APIs need to be called with an authenticated user.
Which solution will meet these requirements with the LEAST amount of effort?

A. Create an internal Application Load Balancer (ALB). Create a target grou
B. Select the Lambda function to cal
C. Use the ALB DNS name to call the API from the VPC.
D. Remove the DNS entry that is associated with the API in API Gatewa
E. Create a hosted zone in Amazon Route 53. Create a CNAME record in the hosted zon
F. Update the API in API Gateway with the CNAME recor
G. Use the CNAME record to call the API from the VPC.
H. Update the API endpoint from Regional to private in API Gatewa
I. Create an interface VPC endpoint in the VP
J. Create a resource policy, and attach it to the AP
K. Use the VPC endpoint to call the API from the VPC.
L. Deploy the Lambda functions inside the VP
M. Provision an EC2 instance, and install an Apache server.From the Apache server, call the Lambda function
N. Use the internal CNAME record of the EC2 instance to call the API from the VPC.

Answer: C

Explanation:
This solution requires the least amount of effort as it only requires to update the API endpoint to private in API Gateway and create an interface VPC endpoint. Then create a resource policy and attach it to the API. This will make the API only accessible from the VPC and still keep the authentication mechanism intact.
Reference:

> https://aws.amazon.com/api-gateway/features/

NEW QUESTION 7
- (Exam Topic 1)
A company uses Amazon S3 to store files and images in a variety of storage classes. The company's S3 costs have increased substantially during the past year. A solutions architect needs to review data trends for the past 12 months and identity the appropriate storage class for the objects.
Which solution will meet these requirements?

A. Download AWS Cost and Usage Reports for the last 12 months of S3 usag
B. Review AWS Trusted Advisor recommendations for cost savings.
C. Use S3 storage class analysi
D. Import data trends into an Amazon QuickSight dashboard to analyze storage trends.
E. Use Amazon S3 Storage Len
F. Upgrade the default dashboard to include advanced metrics for storage trends.
G. Use Access Analyzer for S3. Download the Access Analyzer for S3 report for the last 12 month
H. Import the csvfile to an Amazon QuickSight dashboard.

Answer: B

Explanation:
https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage_lens.html

NEW QUESTION 8
- (Exam Topic 1)
A company is subject to regulatory audits of its financial information. External auditors who use a single AWS account need access to the company's AWS account. A solutions architect must provide the auditors with secure, read-only access to the company's AWS account. The solution must comply with AWS security best practices.
Which solution will meet these requirements?

A. In the company's AWS account, create resource policies for all resources in the account to grant access to the auditors' AWS accoun
B. Assign a unique external ID to the resource policy.
C. In the company's AWS account create an IAM role that trusts the auditors' AWS account Create an IAM policy that has the required permission
D. Attach the policy to the rol
E. Assign a unique external ID to the role's trust policy.
F. In the company's AWS account, create an IAM use
G. Attach the required IAM policies to the IAM user.Create API access keys for the IAM use
H. Share the access keys with the auditors.
I. In the company's AWS account, create an IAM group that has the required permissions Create an IAM user in the company s account for each audito
J. Add the IAM users to the IAM group.

Answer: B

Explanation:
This solution will allow the external auditors to have read-only access to the company's AWS account while being compliant with AWS security best practices. By creating an IAM role, which is a secure and flexible way of granting access to AWS resources, and trusting the auditors' AWS account, the company can ensure that the auditors only have the permissions that are required for their role and nothing more. Assigning a unique external ID to the role's trust policy, it will ensure that only the auditors' AWS account can assume the role.
Reference:
AWS IAM Roles documentation: https://aws.amazon.com/iam/features/roles/ AWS IAM Best practices: https://aws.amazon.com/iam/security-best-practices/

NEW QUESTION 9
- (Exam Topic 1)
A company is using an on-premises Active Directory service for user authentication. The company wants to use the same authentication service to sign in to the company's AWS accounts, which are using AWS Organizations. AWS Site-to-Site VPN connectivity already exists between the on-premises environment and all the company's AWS accounts.
The company's security policy requires conditional access to the accounts based on user groups and roles. User identities must be managed in a single location.
Which solution will meet these requirements?

A. Configure AWS Single Sign-On (AWS SSO) to connect to Active Directory by using SAML 2.0.Enable automatic provisioning by using the System for Cross-domain Identity Management (SCIM) v2.0 protoco
B. Grant access to the AWS accounts by using attribute-based access controls (ABACs).
C. Configure AWS Single Sign-On (AWS SSO) by using AWS SSO as an identity sourc
D. Enable automatic provisioning by using the System for Cross-domain Identity Management (SCIM) v2.0 protoco
E. Grant access to the AWS accounts by using AWS SSO permission sets.
F. In one of the company's AWS accounts, configure AWS Identity and Access Management (IAM) to use a SAML 2.0 identity provide
G. Provision IAM users that are mapped to the federated user
H. Grant access that corresponds to appropriate groups in Active Director
I. Grant access to the required AWS accounts by using cross-account IAM users.
J. In one of the company's AWS accounts, configure AWS Identity and Access Management (IAM) to use an OpenID Connect (OIDC) identity provide
K. Provision IAM roles that grant access to the AWS account for the federated users that correspond to appropriate groups in Active Director
L. Grant access to the required AWS accounts by using cross-account IAM roles.

Answer: D

Explanation:
https://aws.amazon.com/blogs/aws/new-attributes-based-access-control-with-aws-single-sign-on/

NEW QUESTION 10

- (Exam Topic 1)

A company has a latency-sensitive trading platform that uses Amazon DynamoDB as a storage backend. The company configured the DynamoDB table to use on-demand capacity mode. A solutions architect needs to design a solution to improve the performance of the trading platform. The new solution must ensure high availability for the trading platform.

Which solution will meet these requirements with the LEAST latency?

A. Create a two-node DynamoDB Accelerator (DAX) cluster Configure an application to read and write data by using DAX.
B. Create a three-node DynamoDB Accelerator (DAX) cluste
C. Configure an application to read data by using DAX and to write data directly to the DynamoDB table.
D. Create a three-node DynamoDB Accelerator (DAX) cluste
E. Configure an application to read data directly from the DynamoDB table and to write data by using DAX.
F. Create a single-node DynamoD8 Accelerator (DAX) cluste
G. Configure an application to read data by using DAX and to write data directly to the DynamoD8 table.

Answer: B

Explanation:

A DAX cluster can be deployed with one or two nodes for development or test workloads. One- and two-node clusters are not fault-tolerant, and we don't recommend using fewer than three nodes for production use. If a one- or two-node cluster encounters software or hardware errors, the cluster can become unavailable or lose cached data.A DAX cluster can be deployed with one or two nodes for development or test workloads. One and two-node clusters are not fault-tolerant, and we don't recommend using fewer than three nodes for production use. If a one- or two-node cluster encounters software or hardware errors, the cluster can become unavailable or lose cached data.
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DAX.concepts.cluster.html

NEW QUESTION 10

- (Exam Topic 1)

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 rule that runs every hou
F. Configure the rule to invoke one AWS Lambda function that terminates or restores instances from their last backup based on the ta
G. day, and time.

Answer: B

Explanation:

Creating an Amazon EventBridge rule that runs every business day in the evening to stop instances and another rule that runs every business day in the morning to start instances based on the tag will reduce costs with the least operational effort. This approach allows for instances to be stopped during non-business hours when they are not in use, reducing the costs associated with running them. It also allows for instances to be started again in the morning when the development and testing teams need to use them.

NEW QUESTION 14

- (Exam Topic 1)

A company runs a Python script on an Amazon EC2 instance to process data. The script runs every
10 minutes. The script ingests files from an Amazon S3 bucket and processes the files. On average, the script takes approximately 5 minutes to process each file The script will not reprocess a file that the script has already processed.

The company reviewed Amazon CloudWatch metrics and noticed that the EC2 instance is idle for approximately 40% of the time because of the file processing speed. The company wants to make the workload highly available and scalable. The company also wants to reduce long-term management overhead.

Which solution will meet these requirements MOST cost-effectively?

A. Migrate the data processing script to an AWS Lambda functio
B. Use an S3 event notification to invoke the Lambda function to process the objects when the company uploads the objects.
C. Create an Amazon Simple Queue Service (Amazon SQS) queu
D. Configure Amazon S3 to send event notifications to the SQS queu
E. Create an EC2 Auto Scaling group with a minimum size of one instanc
F. Update the data processing script to poll the SQS queue
G. Process the S3 objects that the SQS message identifies.
H. Migrate the data processing script to a container imag
I. Run the data processing container on an EC2 instanc
J. Configure the container to poll the S3 bucket for new objects and to process the resulting objects.
K. Migrate the data processing script to a container image that runs on Amazon Elastic Container Service (Amazon ECS) on AWS Fargat
L. Create an AWS Lambda function that calls the Fargate RunTaskAPI operation when the container processes the fil
M. Use an S3 event notification to invoke the Lambda function.

Answer: D

Explanation:

migrating the data processing script to an AWS Lambda function and using an S3 event notification to invoke the Lambda function to process the objects when the company uploads the objects. This solution meets the company's requirements of high availability and scalability, as well as reducing long-term management

overhead, and is likely to be the most cost-effective option.

NEW QUESTION 17
- (Exam Topic 1)
A company is building a software-as-a-service (SaaS) solution on AWS. The company has deployed an Amazon API Gateway REST API with AWS Lambda integration in multiple AWS Regions and in the same production account.
The company offers tiered pricing that gives customers the ability to pay for the capacity to make a certain number of API calls per second. The premium tier offers up to 3,000 calls per second, and customers are identified by a unique API key. Several premium tier customers in various Regions report that they receive error responses of 429 Too Many Requests from multiple API methods during peak usage hours. Logs indicate that the Lambda function is never invoked.
What could be the cause of the error messages for these customers?

A. The Lambda function reached its concurrency limit.
B. The Lambda function its Region limit for concurrency.
C. The company reached its API Gateway account limit for calls per second.
D. The company reached its API Gateway default per-method limit for calls per second.

Answer: C

Explanation:
https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-request-throttling.html#apig-reques

NEW QUESTION 18
- (Exam Topic 1)
A company built an application based on AWS Lambda deployed in an AWS CloudFormation stack. The last production release of the web application introduced an issue that resulted in an outage lasting several minutes. A solutions architect must adjust the deployment process to support a canary release.
Which solution will meet these requirements?

A. Create an alias for every new deployed version of the Lambda functio
B. Use the AWS CLI update-alias command with the routing-config parameter to distribute the load.
C. Deploy the application into a new CloudFormation stac
D. Use an Amazon Route 53 weighted routing policy to distribute the load.
E. Create a version for every new deployed Lambda functio
F. Use the AWS CLIupdate-function-configuration command with the routing-config parameter to distribute the load.
G. Configure AWS CodeDeploy and use CodeDeployDefault.OneAtATime in the Deployment configuration to distribute the load.

Answer: A

Explanation:
https://aws.amazon.com/blogs/compute/implementing-canary-deployments-of-aws-lambda-functions-with-alias-
https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html

NEW QUESTION 23
- (Exam Topic 1)
A company has introduced a new policy that allows employees to work remotely from their homes if they connect by using a VPN The company Is hosting Internal applications with VPCs in multiple AWS accounts Currently the applications are accessible from the company's on-premises office network through an AWS Site-to-Site VPN connection The VPC in the company's main AWS account has peering connections established with VPCs in other AWS accounts.
A solutions architect must design a scalable AWS Client VPN solution for employees to use while they work from home
What is the MOST cost-effective solution that meets these requirements?

A. Create a Client VPN endpoint in each AWS account Configure required routing that allows access to internal applications
B. Create a Client VPN endpoint in the mam AWS account Configure required routing that allows access to internal applications
C. Create a Client VPN endpoint in the main AWS account Provision a transit gateway that is connected to each AWS account Configure required routing that allows access to internal applications
D. Create a Client VPN endpoint in the mam AWS account Establish connectivity between the Client VPN endpoint and the AWS Site-to-Site VPN

Answer: C

Explanation:
https://docs.aws.amazon.com/vpn/latest/clientvpn-admin/scenario-peered.html

NEW QUESTION 27
- (Exam Topic 1)
A video streaming company recently launched a mobile app for video sharing. The app uploads various files to an Amazon S3 bucket in the us-east-1 Region. The files range in size from 1 GB to 10 GB.
Users who access the app from Australia have experienced uploads that take long periods of time Sometimes the files fail to completely upload for these users . A solutions architect must improve the app' performance for these uploads
Which solutions will meet these requirements? (Select TWO.)

A. Enable S3 Transfer Acceleration on the S3 bucket Configure the app to use the Transfer Acceleration endpoint for uploads
B. Configure an S3 bucket in each Region to receive the upload
C. Use S3 Cross-Region Replication to copy the files to the distribution S3 bucket.
D. Set up Amazon Route 53 with latency-based routing to route the uploads to the nearest S3 bucket Region.
E. Configure the app to break the video files into chunks Use a multipart upload to transfer files to Amazon S3.
F. Modify the app to add random prefixes to the files before uploading

Answer: AD

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/s3-upload-large-files/

Enabling S3 Transfer Acceleration on the S3 bucket and configuring the app to use the Transfer Acceleration endpoint for uploads will improve the app's performance for these uploads by leveraging Amazon CloudFront's globally distributed edge locations to accelerate the uploads. Breaking the video files into chunks and using a multipart upload to transfer files to Amazon S3 will also improve the app's performance by allowing parts of the file to be uploaded in parallel, reducing the overall upload time.

NEW QUESTION 28
- (Exam Topic 1)
A company is building an electronic document management system in which users upload their documents. The application stack is entirely serverless and runs on AWS in the eu-central-1 Region. The system includes a web application that uses an Amazon CloudFront distribution for delivery with Amazon S3 as the origin. The web application communicates with Amazon API Gateway Regional endpoints. The API Gateway APIs call AWS Lambda functions that store metadata in an Amazon Aurora Serverless database and put the documents into an S3 bucket.
The company is growing steadily and has completed a proof of concept with its largest customer. The company must improve latency outside of Europe.
Which combination of actions will meet these requirements? (Select TWO.)

A. Enable S3 Transfer Acceleration on the S3 bucke
B. Ensure that the web application uses the Transfer Acceleration signed URLs.
C. Create an accelerator in AWS Global Accelerato
D. Attach the accelerator to the CloudFront distribution.
E. Change the API Gateway Regional endpoints to edge-optimized endpoints.
F. Provision the entire stack in two other locations that are spread across the worl
G. Use global databases on the Aurora Serverless cluster.
H. Add an Amazon RDS proxy between the Lambda functions and the Aurora Serverless database.

Answer: AC

Explanation:
https://aws.amazon.com/global-accelerator/faqs/

NEW QUESTION 32
- (Exam Topic 1)
A company needs to implement a patching process for its servers. The on-premises servers and Amazon EC2 instances use a variety of tools to perform patching. Management requires a single report showing the patch status of all the servers and instances.
Which set of actions should a solutions architect take to meet these requirements?

A. Use AWS Systems Manager to manage patches on the on-premises servers and EC2 instance
B. Use Systems Manager to generate patch compliance reports.
C. Use AWS OpsWorks to manage patches on the on-premises servers and EC2 instance
D. Use Amazon OuickSight integration with OpsWorks to generate patch compliance reports.
E. Use an Amazon EventBridge (Amazon CloudWatch Events) rule to apply patches by scheduling an AWS Systems Manager patch remediation jo
F. Use Amazon Inspector to generate patch compliance reports.
G. Use AWS OpsWorks to manage patches on the on-premises servers and EC2 instance
H. Use AWS X-Ray to post the patch status to AWS Systems Manager OpsCenter to generate patch compliance reports.

Answer: A

Explanation:
https://docs.aws.amazon.com/systems-manager/latest/userguide/systems-manager-patch.html

NEW QUESTION 34
- (Exam Topic 1)
A company is planning to store a large number of archived documents and make the documents available to employees through the corporate intranet. Employees will access the system by connecting through a client VPN service that is attached to a VPC. The data must not be accessible to the public.
The documents that the company is storing are copies of data that is held on physical media elsewhere. The number of requests will be low. Availability and speed of retrieval are not concerns of the company.
Which solution will meet these requirements at the LOWEST cost?

A. Create an Amazon S3 bucke
B. Configure the S3 bucket to use the S3 One Zone-Infrequent Access (S3 One Zone-IA) storage class as defaul
C. Configure the S3 bucket for website hostin
D. Create an S3 interface endpoin
E. Configure the S3 bucket to allow access only through that endpoint.
F. Launch an Amazon EC2 instance that runs a web serve
G. Attach an Amazon Elastic File System (Amazon EFS) file system to store the archived data in the EFS One Zone-Infrequent Access (EFS One Zone-IA) storage class Configure the instance security groups to allow access only from private networks.
H. Launch an Amazon EC2 instance that runs a web server Attach an Amazon Elastic Block Store (Amazon EBS) volume to store the archived dat
I. Use the Cold HDD (sc1) volume typ
J. Configure the instance security groups to allow access only from private networks.
K. Create an Amazon S3 bucke
L. Configure the S3 bucket to use the S3 Glacier Deep Archive storage class as defaul
M. Configure the S3 bucket for website hostin
N. Create an S3 interface endpoin
O. Configure the S3 bucket to allow access only through that endpoint.

Answer: D

Explanation:
The S3 Glacier Deep Archive storage class is the lowest-cost storage class offered by Amazon S3, and it is designed for archival data that is accessed infrequently and for which retrieval time of several hours is acceptable. S3 interface endpoint for the VPC ensures that access to the bucket is only from resources within the VPC and this will meet the requirement of not being accessible to the public. And also, S3 bucket can be configured for website hosting, and this will allow employees to access the documents through the corporate intranet. Using an EC2 instance and a file system or block store would be more expensive and

unnecessary because the number of requests to the data will be low and availability and speed of retrieval are not concerns. Additionally, using Amazon S3 bucket will provide durability, scalability and availability of data.

NEW QUESTION 39
- (Exam Topic 1)
A company has registered 10 new domain names. The company uses the domains for online marketing. The company needs a solution that will redirect online visitors to a specific URL for each domain. All domains and target URLs are defined in a JSON document. All DNS records are managed by Amazon Route 53.
A solutions architect must implement a redirect service that accepts HTTP and HTTPS requests.
Which combination of steps should the solutions architect take to meet these requirements with the LEAST amount of operational effort? (Choose three.)

A. Create a dynamic webpage that runs on an Amazon EC2 instanc
B. Configure the webpage to use the JSON document in combination with the event message to look up and respond with a redirect URL.
C. Create an Application Load Balancer that includes HTTP and HTTPS listeners.
D. Create an AWS Lambda function that uses the JSON document in combination with the event message to look up and respond with a redirect URL.
E. Use an Amazon API Gateway API with a custom domain to publish an AWS Lambda function.
F. Create an Amazon CloudFront distributio
G. Deploy a Lambda@Edge function.
H. Create an SSL certificate by using AWS Certificate Manager (ACM). Include the domains as Subject Alternative Names.

Answer: CEF

Explanation:
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-edge-how-it-works-tutorial.ht

NEW QUESTION 40
- (Exam Topic 1)
A global media company is planning a multi-Region deployment of an application. Amazon DynamoDB global tables will back the deployment to keep the user experience consistent across the two continents where users are concentrated. Each deployment will have a public Application Load Balancer (ALB). The company manages public DNS internally. The company wants to make the application available through an apex domain.
Which solution will meet these requirements with the LEAST effort?

A. Migrate public DNS to Amazon Route 53. Create CNAME records for the apex domain to point to the AL
B. Use a geolocation routing policy to route traffic based on user location.
C. Place a Network Load Balancer (NLB) in front of the AL
D. Migrate public DNS to Amazon Route 53.Create a CNAME record for the apex domain to point to the NLB's static IP addres
E. Use a geolocation routing policy to route traffic based on user location.
F. Create an AWS Global Accelerator accelerator with multiple endpoint groups that target endpoints in appropriate AWS Region
G. Use the accelerator's static IP address to create a record in public DNS for the apex domain.
H. Create an Amazon API Gateway API that is backed by AWS Lambda in one of the AWS Regions.Configure a Lambda function to route traffic to application deployments by using the round robin metho
I. Create CNAME records for the apex domain to point to the API's URL.

Answer: C

Explanation:
AWS Global Accelerator is a service that directs traffic to optimal endpoints (in this case, the Application Load Balancer) based on the health of the endpoints and network routing. It allows you to create an accelerator that directs traffic to multiple endpoint groups, one for each Region where the application is deployed. The accelerator uses the AWS global network to optimize the traffic routing to the healthy endpoint.
By using Global Accelerator, the company can use a single static IP address for the apex domain, and traffic will be directed to the optimal endpoint based on the user's location, without the need for additional load balancers or routing policies.
Reference:
AWS Global Accelerator documentation: https://aws.amazon.com/global-accelerator/ Routing User Traffic to the Optimal AWS Region using Global Accelerator documentation:
https://aws.amazon.com/blogs/networking-and-content-delivery/routing-user-traffic-to-the-optimal-aws-region-u

NEW QUESTION 45
- (Exam Topic 1)
A company is storing data in several Amazon DynamoDB tables. A solutions architect must use a serverless architecture to make the data accessible publicly through a simple API over HTTPS. The solution must scale automatically in response to demand.
Which solutions meet these requirements? (Choose two.)

A. Create an Amazon API Gateway REST AP
B. Configure this API with direct integrations to DynamoDB by using API Gateway's AWS integration type.
C. Create an Amazon API Gateway HTTP AP
D. Configure this API with direct integrations to Dynamo DB by using API Gateway's AWS integration type.
E. Create an Amazon API Gateway HTTP AP
F. Configure this API with integrations to AWS Lambda functions that return data from the DynamoDB tables.
G. Create an accelerator in AWS Global Accelerato
H. Configure this accelerator with AWS Lambda@Edge function integrations that return data from the DynamoDB tables.
I. Create a Network Load Balance
J. Configure listener rules to forward requests to the appropriate AWS Lambda functions

Answer: AC

Explanation:
https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-overview-developer-experience.htm

NEW QUESTION 49
- (Exam Topic 1)

A video processing company has an application that downloads images from an Amazon S3 bucket, processes the images, stores a transformed image in a second S3 bucket, and updates metadata about the image in an Amazon DynamoDB table. The application is written in Node.js and runs by using an AWS Lambda function. The Lambda function is invoked when a new image is uploaded to Amazon S3.

The application ran without incident for a while. However, the size of the images has grown significantly. The Lambda function is now failing frequently with timeout errors. The function timeout is set to its maximum value. A solutions architect needs to refactor the application's architecture to prevent invocation failures. The company does not want to manage the underlying infrastructure.

Which combination of steps should the solutions architect take to meet these requirements? (Choose two.)

A. Modify the application deployment by building a Docker image that contains the application code.Publish the image to Amazon Elastic Container Registry (Amazon ECR).

B. Create a new Amazon Elastic Container Service (Amazon ECS) task definition with a compatibility type of AWS Fargat

C. Configure the task definition to use the new image in Amazon Elastic Container Registry (Amazon ECR). Adjust the Lambda function to invoke an ECS task by using the ECS task definition when a new file arrives in Amazon S3.

D. Create an AWS Step Functions state machine with a Parallel state to invoke the Lambda function.Increase the provisioned concurrency of the Lambda function.

E. Create a new Amazon Elastic Container Service (Amazon ECS) task definition with a compatibility type of Amazon EC2. Configure the task definition to use the new image in Amazon Elastic Container Registry (Amazon ECR). Adjust the Lambda function to invoke an ECS task by using the ECS task definition when a new file arrives in Amazon S3.

F. Modify the application to store images on Amazon Elastic File System (Amazon EFS) and to store metadata on an Amazon RDS DB instanc

G. Adjust the Lambda function to mount the EFS file share.

Answer: AB

Explanation:
A. Modify the application deployment by building a Docker image that contains the application code. Publish the image to Amazon Elastic Container Registry (Amazon ECR). - This step is necessary to package the application code in a container and make it available for running on ECS. B. Create a new Amazon Elastic Container Service (Amazon ECS) task definition with a compatibility type of AWS Fargate. Configure the task definition to use the new image in Amazon Elastic Container Registry (Amazon ECR). Adjust the Lambda function to invoke an ECS task by using the ECS task definition when a new file arrives in Amazon S3.

NEW QUESTION 51
- (Exam Topic 1)
A company wants to use AWS to create a business continuity solution in case the company's main on-premises application fails. The application runs on physical servers that also run other applications. The on-premises application that the company is planning to migrate uses a MySQL database as a data store. All the company's on-premises applications use operating systems that are compatible with Amazon EC2.

Which solution will achieve the company's goal with the LEAST operational overhead?

A. Install the AWS Replication Agent on the source servers, including the MySQL server

B. Set up replication for all server

C. Launch test instances for regular drill

D. Cut over to the test instances to fail over the workload in the case of a failure event.

E. Install the AWS Replication Agent on the source servers, including the MySQL server

F. Initialize AWS Elastic Disaster Recovery in the target AWS Regio

G. Define the launch setting

H. Frequently perform failover and fallback from the most recent point in time.

I. Create AWS Database Migration Service (AWS DMS) replication servers and a target Amazon Aurora MySQL DB cluster to host the databas

J. Create a DMS replication task to copy the existing data to the target DB cluste

K. Create a local AWS Schema Conversion Tool (AWS SCT) change data capture (CDC) task to keep the data synchronize

L. Install the rest of the software on EC2 instances by starting with a compatible base AMI.

M. Deploy an AWS Storage Gateway Volume Gateway on premise

N. Mount volumes on all on-premises server

O. Install the application and the MySQL database on the new volume

P. Take regular snapshot

Q. Install all the software on EC2 Instances by starting with a compatible base AM

R. Launch a Volume Gateway on an EC2 instanc

S. Restore the volumes from the latest snapsho

T. Mount the new volumes on the EC2 instances in the case of a failure event.

Answer: B

Explanation:
https://docs.aws.amazon.com/drs/latest/userguide/what-is-drs.html https://docs.aws.amazon.com/drs/latest/userguide/recovery-workflow-gs.html

NEW QUESTION 53
- (Exam Topic 1)
A company is running an application in the AWS Cloud. Recent application metrics show inconsistent
response times and a significant increase in error rates. Calls to third-party services are causing the delays. Currently, the application calls third-party services synchronously by directly invoking an AWS Lambda function.

A solutions architect needs to decouple the third-party service calls and ensure that all the calls are eventually completed.

Which solution will meet these requirements?

A. Use an Amazon Simple Queue Service (Amazon SQS) queue to store events and invoke the Lambda function.

B. Use an AWS Step Functions state machine to pass events to the Lambda function.

C. Use an Amazon EventBridge rule to pass events to the Lambda function.

D. Use an Amazon Simple Notification Service (Amazon SNS) topic to store events and Invoke the Lambda function.

Answer: A

Explanation:
Using an SQS queue to store events and invoke the Lambda function will decouple the third-party service calls and ensure that all the calls are eventually completed. SQS allows you to store messages in a queue and process them asynchronously, which eliminates the need for the application to wait for a response from the third-party service. The messages will be stored in the SQS queue until they are processed by the Lambda function, even if the Lambda function is currently unavailable or busy. This will ensure that all the calls are eventually completed, even if there are delays or errors.

AWS Step Functions state machines can also be used to pass events to the Lambda function, but it would require additional management and configuration to set up the state machine, which would increase operational overhead.

Amazon EventBridge rule can also be used to pass events to the Lambda function, but it would not provide the same level of decoupling and reliability as SQS.

Using Amazon Simple Notification Service (Amazon SNS) topic to store events and Invoke the Lambda function, is similar to SQS, but SNS is a publish-subscribe messaging service and SQS is a queue service. SNS is used for sending messages to multiple recipients, SQS is used for sending messages to a single recipient, so SQS is more appropriate for this use case.

References:

- AWS SQS
- AWS Step Functions
- AWS EventBridge
- AWS SNS

NEW QUESTION 55
- (Exam Topic 1)
A solutions architect is auditing the security setup of an AWS Lambda function for a company. The Lambda function retrieves the latest changes from an Amazon Aurora database. The Lambda function and the database run in the same VPC. Lambda environment variables are providing the database credentials to the Lambda function.

The Lambda function aggregates data and makes the data available in an Amazon S3 bucket that is configured for server-side encryption with AWS KMS managed encryption keys (SSE-KMS). The data must not travel across the internet. If any database credentials become compromised, the company needs a solution that minimizes the impact of the compromise.

What should the solutions architect recommend to meet these requirements?

A. Enable IAM database authentication on the Aurora DB cluste
B. Change the IAM role for the Lambda function to allow the function to access the database by using IAM database authenticatio
C. Deploy a gateway VPC endpoint for Amazon S3 in the VPC.
D. Enable IAM database authentication on the Aurora DB cluste
E. Change the IAM role for the Lambda function to allow the function to access the database by using IAM database authenticatio
F. Enforce HTTPS on the connection to Amazon S3 during data transfers.
G. Save the database credentials in AWS Systems Manager Parameter Stor
H. Set up password rotation on the credentials in Parameter Stor
I. Change the IAM role for the Lambda function to allow the function to access Parameter Stor
J. Modify the Lambda function to retrieve the credentials from Parameter Stor
K. Deploy a gateway VPC endpoint for Amazon S3 in the VPC.
L. Save the database credentials in AWS Secrets Manage
M. Set up password rotation on the credentials in Secrets Manage
N. Change the IAM role for the Lambda function to allow the function to access Secrets Manage
O. Modify the Lambda function to retrieve the credentials Om Secrets Manage
P. Enforce HTTPS on the connection to Amazon S3 during data transfers.

Answer: A

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/UsingWithRDS.IAMDBAuth.html

NEW QUESTION 56
- (Exam Topic 1)
A company is hosting an image-processing service on AWS in a VPC. The VPC extends across two Availability Zones. Each Availability Zone contains one public subnet and one private subnet.

The service runs on Amazon EC2 instances in the private subnets. An Application Load Balancer in the public subnets is in front of the service. The service needs to communicate with the internet and does so through two NAT gateways. The service uses Amazon S3 for image storage. The EC2 instances retrieve approximately 1 ¢' of data from an S3 bucket each day.

The company has promoted the service as highly secure. A solutions architect must reduce cloud expenditures as much as possible without compromising the service's security posture or increasing the time spent on ongoing operations.

Which solution will meet these requirements?

A. Replace the NAT gateways with NAT instance
B. In the VPC route table, create a route from the private subnets to the NAT instances.
C. Move the EC2 instances to the public subnet
D. Remove the NAT gateways.
E. Set up an S3 gateway VPC endpoint in the VP
F. Attach an endpoint policy to the endpoint to allow the required actions on the S3 bucket.
G. Attach an Amazon Elastic File System (Amazon EFS) volume to the EC2 instance
H. Host the image on the EFS volume.

Answer: C

Explanation:
Create Amazon S3 gateway endpoint in the VPC and add a VPC endpoint policy. This VPC endpoint policy will have a statement that allows S3 access only via access points owned by the organization.

NEW QUESTION 57
- (Exam Topic 1)
A start up company hosts a fleet of Amazon EC2 instances in private subnets using the latest Amazon Linux 2 AMI. The company's engineers rely heavily on SSH access to the instances for troubleshooting.

The company's existing architecture includes the following:

• A VPC with private and public subnets, and a NAT gateway
• Site-to-Site VPN for connectivity with the on-premises environment
• EC2 security groups with direct SSH access from the on-premises environment

The company needs to increase security controls around SSH access and provide auditing of commands executed by the engineers.
Which strategy should a solutions architect use?

A. Install and configure EC2 Instance Connect on the fleet of EC2 instance
B. Remove all security group rules attached to EC2 instances that allow inbound TCP on port 22. Advise the engineers to remotely access the instances by using the EC2 Instance Connect CLI.
C. Update the EC2 security groups to only allow inbound TCP on port 22 to the IP addresses of the engineer's device
D. Install the Amazon CloudWatch agent on all EC2 instances and send operating system audit logs to CloudWatch Logs.
E. Update the EC2 security groups to only allow inbound TCP on port 22 to the IP addresses of the engineer's device
F. Enable AWS Config for EC2 security group resource change
G. Enable AWS Firewall Manager and apply a security group policy that automatically remediates changes to rules.
H. Create an IAM role with the AmazonSSMManagedInstanceCore managed policy attache
I. Attach the IAM role to all the EC2 instance
J. Remove all security group rules attached to the EC2 instances that allow inbound TCP on port 22. Have the engineers install the AWS Systems Manager Session Manager plugin for their devices and remotely access the instances by using the start-session API call from Systems Manager.

Answer: D

Explanation:
Allows client machines to be able to connect to Session Manager using the AWS CLI instead of going through the AWS EC2 or AWS Server Manager console. https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager-working-with-install-plugin.ht https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager-working-with-install-plugin.ht

NEW QUESTION 60
- (Exam Topic 1)
A company wants to use a third-party software-as-a-service (SaaS) application. The third-party SaaS application is consumed through several API calls. The third-party SaaS application also runs on AWS inside a VPC.
The company will consume the third-party SaaS application from inside a VPC. The company has internal security policies that mandate the use of private connectivity that does not traverse the internet. No resources that run in the company VPC are allowed to be accessed from outside the company's VPC. All permissions must conform to the principles of least privilege.
Which solution meets these requirements?

A. Create an AWS PrivateLink interface VPC endpoin
B. Connect this endpoint to the endpoint service that the third-party SaaS application provide
C. Create a security group to limit the access to the endpoin
D. Associate the security group with the endpoint.
E. Create an AWS Site-to-Site VPN connection between the third-party SaaS application and the company VP
F. Configure network ACLs to limit access across the VPN tunnels.
G. Create a VPC peering connection between the third-party SaaS application and the company VPUpdate route tables by adding the needed routes for the peering connection.
H. Create an AWS PrivateLink endpoint servic
I. Ask the third-party SaaS provider to create an interface VPC endpoint for this endpoint servic
J. Grant permissions for the endpoint service to the specific account of the third-party SaaS provider.

Answer: A

Explanation:
Reference architecture - https://docs.aws.amazon.com/vpc/latest/privatelink/privatelink-access-saas.html Note from documentation that Interface Endpoint is at client side

NEW QUESTION 64
- (Exam Topic 2)
A company needs to audit the security posture of a newly acquired AWS account. The company's data security team requires a notification only when an Amazon S3 bucket becomes publicly exposed. The company has already established an Amazon Simple Notification Service (Amazon SNS) topic that has the data security team's email address subscribed.
Which solution will meet these requirements?

A. Create an S3 event notification on all S3 buckets for the isPublic even
B. Select the SNS topic as the target for the event notifications.
C. Create an analyzer in AWS Identity and Access Management Access Analyze
D. Create an Amazon EventBridge rule for the event type "Access Analyzer Finding" with a filter for "isPublic: true." Select the SNS topic as the EventBridge rule target.
E. Create an Amazon EventBridge rule for the event type "Bucket-Level API Call via CloudTrail" with a filter for "PutBucketPolicy." Select the SNS topic as the EventBridge rule target.
F. Activate AWS Config and add the cloudtrail-s3-dataevents-enabled rul
G. Create an Amazon EventBridge rule for the event type "Config Rules Re-evaluation Status" with a filter for "NON_COMPLIANT." Select the SNS topic as the EventBridge rule target.

Answer: B

Explanation:
Access Analyzer is to assess the access policy. https://docs.aws.amazon.com/ja_jp/AmazonS3/latest/userguide/access-control-block-public-access.html

NEW QUESTION 65
- (Exam Topic 2)
A company is running a compute workload by using Amazon EC2 Spot Instances that are in an Auto Scaling group. The launch template uses two placement groups and a single instance type.
Recently, a monitoring system reported Auto Scaling instance launch failures that correlated with longer wait times for system users. The company needs to improve the overall reliability of the workload.
Which solution will meet this requirement?

A. Replace the launch template with a launch configuration to use an Auto Scaling group that uses attribute-based instance type selection.
B. Create a new launch template version that uses attribute-based instance type selectio
C. Configure the Auto Scaling group to use the new launch template version.
D. Update the launch template Auto Scaling group to increase the number of placement groups.
E. Update the launch template to use a larger instance type.

Answer: B

Explanation:
https://docs.aws.amazon.com/autoscaling/ec2/userguide/create-asg-instance-type-requirements.html#use-attribut

NEW QUESTION 70
- (Exam Topic 2)
A company processes environment data. The has a set up sensors to provide a continuous stream of data from different areas in a city. The data is available in JSON format.
The company wants to use an AWS solution to send the data to a database that does not require fixed schemas for storage. The data must be send in real time.
Which solution will meet these requirements?

A. Use Amazon Kinesis Data Firehouse to send the data to Amazon Redshift.
B. Use Amazon Kinesis Data streams to send the data to Amazon DynamoDB.
C. Use Amazon Managed Streaming for Apache Kafka (Amazon MSK) to send the data to Amazon Aurora.
D. Use Amazon Kinesis Data firehouse to send the data to Amazon Keyspaces (for Apache Cassandra).

Answer: B

Explanation:
Amazon Kinesis Data Streams is a service that enables real-time data ingestion and processing. Amazon DynamoDB is a NoSQL database that does not require fixed schemas for storage. By using Kinesis Data Streams and DynamoDB, the company can send the JSON data to a database that can handle schemaless data in real time. References:
≫ https://docs.aws.amazon.com/streams/latest/dev/introduction.html
≫ https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html

NEW QUESTION 73
- (Exam Topic 2)
A telecommunications company is running an application on AWS. The company has set up an AWS Direct Connect connection between the company's on-premises data center and AWS. The company deployed the application on Amazon EC2 instances in multiple Availability Zones behind an internal Application Load Balancer (ALB). The company's clients connect from the on-premises network by using HTTPS. The TLS terminates in the ALB. The company has multiple target groups and uses path-based routing to forward requests based on the URL path.
The company is planning to deploy an on-premises firewall appliance with an allow list that is based on IP address. A solutions architect must develop a solution to allow traffic flow to AWS from the on-premises network so that the clients can continue to access the application.
Which solution will meet these requirements?

A. Configure the existing ALB to use static IP addresse
B. Assign IP addresses in multiple Availability Zones to the AL
C. Add the ALB IP addresses to the firewall appliance.
D. Create a Network Load Balancer (NLB). Associate the NLB with one static IP addresses in multiple Availability Zone
E. Create an ALB-type target group for the NLB and add the existing ALAdd the NLB IP addresses to the firewall applianc
F. Update the clients to connect to the NLB.
G. Create a Network Load Balancer (NLB). Associate the LNB with one static IP addresses in multiple Availability Zone
H. Add the existing target groups to the NL
I. Update the clients to connect to the NL
J. Delete the ALB Add the NLB IP addresses to the firewall appliance.
K. Create a Gateway Load Balancer (GWLB). Assign static IP addresses to the GWLB in multiple Availability Zone
L. Create an ALB-type target group for the GWLB and add the existing AL
M. Add the GWLB IP addresses to the firewall applianc
N. Update the clients to connect to the GWLB.

Answer: B

Explanation:
The company should create a Network Load Balancer (NLB) and associate it with one static IP address in multiple Availability Zones. The company should also create an ALB-type target group for the NLB and add the existing ALB. The company should add the NLB IP addresses to the firewall appliance and update the clients to connect to the NLB. This solution will allow traffic flow to AWS from the on-premises network by using static IP addresses that can be added to the firewall appliance's allow list. The NLB will forward requests to the ALB, which will use path-based routing to forward requests to the target groups.

NEW QUESTION 78
- (Exam Topic 2)
A company has built a high performance computing (HPC) cluster in AWS tor a tightly coupled workload that generates a large number of shared files stored in Amazon EFS. The cluster was performing well when the number of Amazon EC2 instances in the cluster was 100. However, when the company increased the cluster size to 1,000 EC2 instances, overall performance was well below expectations.
Which collection of design choices should a solutions architect make to achieve the maximum performance from the HPC cluster? (Select THREE.)

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

Answer: ACF

Explanation:
* A. High performance computing (HPC) workload cluster should be in a single AZ.
* C. Elastic Fabric Adapter (EFA) is a network device that you can attach to your Amazon EC2 instances to accelerate High Performance Computing (HPC)
* F. Amazon FSx for Lustre - Use it for workloads where speed matters, such as machine learning, high performance computing (HPC), video processing, and financial modeling.
Cluster – packs instances close together inside an Availability Zone. This strategy enables workloads to achieve the low-latency network performance necessary for tightly-coupled node-to-node communication that is typical of HPC applications.
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html

NEW QUESTION 82
- (Exam Topic 2)
A company manufactures smart vehicles. The company uses a custom application to collect vehicle data. The vehicles use the MQTT protocol to connect to the application.
The company processes the data in 5-minute intervals. The company then copies vehicle telematics data to on-premises storage. Custom applications analyze this data to detect anomalies.
The number of vehicles that send data grows constantly. Newer vehicles generate high volumes of data. The on-premises storage solution is not able to scale for peak traffic, which results in data loss. The company must modernize the solution and migrate the solution to AWS to resolve the scaling challenges.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use AWS IOT Greengrass to send the vehicle data to Amazon Managed Streaming for Apache Kafka (Amazon MSK). Create an Apache Kafka application to store the data in Amazon S3. Use a pretrained model in Amazon SageMaker to detect anomalies.
B. Use AWS IOT Core to receive the vehicle dat
C. Configure rules to route data to an Amazon Kinesis Data Firehose delivery stream that stores the data in Amazon S3. Create an Amazon Kinesis Data Analytics application that reads from the delivery stream to detect anomalies.
D. Use AWS IOT FleetWise to collect the vehicle dat
E. Send the data to an Amazon Kinesis data stream.Use an Amazon Kinesis Data Firehose delivery stream to store the data in Amazon S3. Use the built-in machine learning transforms in AWS Glue to detect anomalies.
F. Use Amazon MQ for RabbitMQ to collect the vehicle dat
G. Send the data to an Amazon Kinesis Data Firehose delivery stream to store the data in Amazon S3. Use Amazon Lookout for Metrics to detect anomalies.

Answer: B

Explanation:
Using AWS IoT Core to receive the vehicle data will enable connecting the smart vehicles to the cloud using the MQTT protocol1. AWS IoT Core is a platform that enables you to connect devices to AWS Services and other devices, secure data and interactions, process and act upon device data, and enable applications to interact with devices even when they are offline2. Configuring rules to route data to an Amazon Kinesis Data Firehose delivery stream that stores the data in Amazon S3 will enable processing and storing the vehicle data in a scalable and reliable way3. Amazon Kinesis Data Firehose is a fully managed service that delivers real-time streaming data to destinations such as Amazon S3. Creating an Amazon Kinesis Data Analytics application that reads from the delivery stream to detect anomalies will enable analyzing the vehicle data using SQL queries or Apache Flink applications. Amazon Kinesis Data Analytics is a fully managed service that enables you to process and analyze streaming data using SQL or Java.

NEW QUESTION 87
- (Exam Topic 2)
A company operates an on-premises software-as-a-service (SaaS) solution that ingests several files daily. The company provides multiple public SFTP endpoints to its customers to facilitate the file transfers. The customers add the SFTP endpoint IP addresses to their firewall allow list for outbound traffic. Changes to the SFTP endmost IP addresses are not permitted.
The company wants to migrate the SaaS solution to AWS and decrease the operational overhead of the file transfer service.
Which solution meets these requirements?

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

Answer: A

Explanation:
Bring your own IP addresses (BYOIP) You can bring part or all of your publicly routable IPv4 or IPv6 address range from your on-premises network to your AWS account. You continue to own the address range, but AWS advertises it on the internet by default. After you bring the address range to AWS, it appears in your AWS account as an address pool. https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-byoip.html AWS Transfer for SFTP enables you to easily move your file transfer workloads that use the Secure Shell File Transfer Protocol (SFTP) to AWS without needing to modify your applications or manage any SFTP servers. https://aws.amazon.com/about-aws/whats-new/2018/11/aws-transfer-for-sftp-fully-managed-sftp-for-s3/

NEW QUESTION 90
- (Exam Topic 2)
A company needs to architect a hybrid DNS solution. This solution will use an Amazon Route 53 private hosted zone for the domain cloud.example.com for the resources stored within VPCs.
The company has the following DNS resolution requirements:
• On-premises systems should be able to resolve and connect to cloud.example.com.
• All VPCs should be able to resolve cloud.example.com.

There is already an AWS Direct Connect connection between the on-premises corporate network and AWS Transit Gateway. Which architecture should the company use to meet these requirements with the HIGHEST performance?

A. Associate the private hosted zone to all the VPC
B. Create a Route 53 inbound resolver in the shared services VP
C. Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver.
D. Associate the private hosted zone to all the VPC
E. Deploy an Amazon EC2 conditional forwarder in the shared services VP
F. Attach all VPCs to the transit gateway and create forwarding rules in theon-premises DNS server for cloud.example.com that point to the conditional forwarder.
G. Associate the private hosted zone to the shared services VP
H. Create a Route 53 outbound resolver in the shared services VP
I. Attach all VPCs to the transit gateway and create forwarding rules in theon-premises DNS server for cloud.example.com that point to the outbound resolver.
J. Associate the private hosted zone to the shared services VP
K. Create a Route 53 inbound resolver in the shared services VP
L. Attach the shared services VPC to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver.

Answer: A

Explanation:
Amazon Route 53 Resolver is a managed DNS resolver service from Route 53 that helps to create conditional forwarding rules to redirect query traffic1. By associating the private hosted zone to all the VPCs, the solutions architect can enable DNS resolution for cloud.example.com within the VPCs. By creating a Route 53 inbound resolver in the shared services VPC, the solutions architect can enable DNS resolution for cloud.example.com from on-premises systems. By attaching all VPCs to the transit gateway, the solutions architect can enable connectivity between the VPCs and the on-premises network through AWS Direct Connect. By creating forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver, the solutions architect can direct DNS queries for cloud.example.com to the Route 53 Resolver endpoint in AWS. This solution will provide the highest performance as it leverages Route 53 Resolver's optimized routing and caching capabilities.
References: 1: https://aws.amazon.com/route53/resolver/

NEW QUESTION 94
- (Exam Topic 2)
A company uses a Grafana data visualization solution that runs on a single Amazon EC2 instance to monitor the health of the company's AWS workloads. The company has invested time and effort to create dashboards that the company wants to preserve. The dashboards need to be highly available and cannot be down for longer than 10 minutes. The company needs to minimize ongoing maintenance.
Which solution will meet these requirements with the LEAST operational overhead?

A. Migrate to Amazon CloudWatch dashboard
B. Recreate the dashboards to match the existing Grafana dashboard
C. Use automatic dashboards where possible.
D. Create an Amazon Managed Grafana workspac
E. Configure a new Amazon CloudWatch data source.Export dashboards from the existing Grafana instanc
F. Import the dashboards into the new workspace.
G. Create an AMI that has Grafana pre-installe
H. Store the existing dashboards in Amazon Elastic File System (Amazon EFS). Create an Auto Scaling group that uses the new AM
I. Set the Auto Scaling group's minimum, desired, and maximum number of instances to on
J. Create an Application Load Balancer that serves at least two Availability Zones.
K. Configure AWS Backup to back up the EC2 instance that runs Grafana once each hou
L. Restore the EC2 instance from the most recent snapshot in an alternate Availability Zone when required.

Answer: C

Explanation:
By creating an AMI that has Grafana pre-installed and storing the existing dashboards in Amazon Elastic File System (Amazon EFS) it allows for faster and more efficient scaling, and by creating an Auto Scaling group that uses the new AMI and setting the Auto Scaling group's minimum, desired, and maximum number of instances to one and creating an Application Load Balancer that serves at least two Availability Zones, it ensures high availability and minimized downtime.

NEW QUESTION 97
- (Exam Topic 2)
A company is running a web application in a VPC. The web application runs on a group of Amazon EC2 instances behind an Application Load Balancer (ALB). The ALB is using AWS WAF.
An external customer needs to connect to the web application. The company must provide IP addresses to all external customers.
Which solution will meet these requirements with the LEAST operational overhead?

A. Replace the ALB with a Network Load Balancer (NLB). Assign an Elastic IP address to the NLB.
B. Allocate an Elastic IP addres
C. Assign the Elastic IP address to the ALProvide the Elastic IP address to the customer.
D. Create an AWS Global Accelerator standard accelerato
E. Specify the ALB as the accelerator's endpoint.Provide the accelerator's IP addresses to the customer.
F. Configure an Amazon CloudFront distributio
G. Set the ALB as the origi
H. Ping the distribution's DNS name to determine the distribution's public IP addres
I. Provide the IP address to the customer.

Answer: C

Explanation:
https://docs.aws.amazon.com/global-accelerator/latest/dg/about-accelerators.alb-accelerator.html Option A is wrong. AWS WAF does not support associating with NLB.
https://docs.aws.amazon.com/waf/latest/developerguide/waf-chapter.html Option B is wrong. An ALB does not support an Elastic IP address.
https://aws.amazon.com/elasticloadbalancing/features/

NEW QUESTION 98
- (Exam Topic 2)
A company runs an application in an on-premises data center. The application gives users the ability to upload media files. The files persist in a file server. The web application has many users. The application server is overutilized, which causes data uploads to fail occasionally. The company frequently adds new storage to the file server. The company wants to resolve these challenges by migrating the application to AWS.
Users from across the United States and Canada access the application. Only authenticated users should have the ability to access the application to upload files. The company will consider a solution that refactors the application, and the company needs to accelerate application development.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use AWS Application Migration Service to migrate the application server to Amazon EC2 instances.Create an Auto Scaling group for the EC2 instance
B. Use an Application Load Balancer to distribute the request
C. Modify the application to use Amazon S3 to persist the file
D. Use Amazon Cognito to authenticate users.
E. Use AWS Application Migration Service to migrate the application server to Amazon EC2 instances.Create an Auto Scaling group for the EC2 instance
F. Use an Application Load Balancer to distribute the request
G. Set up AWS IAM Identity Center (AWS Single Sign-On) to give users the ability to sign in to the applicatio
H. Modify the application to use Amazon S3 to persist the files.
I. Create a static website for uploads of media file
J. Store the static assets in Amazon S3. Use AWS AppSync to create an AP
K. Use AWS Lambda resolvers to upload the media files to Amazon S3. Use Amazon Cognito to authenticate users.
L. Use AWS Amplify to create a static website for uploads of media file
M. Use Amplify Hosting to serve the website through Amazon CloudFron
N. Use Amazon S3 to store the uploaded media file
O. Use Amazon Cognito to authenticate users.

Answer: D

Explanation:
The company should use AWS Amplify to create a static website for uploads of media files. The company should use Amplify Hosting to serve the website through Amazon CloudFront. The company should use Amazon S3 to store the uploaded media files. The company should use Amazon Cognito to authenticate users. This solution will meet the requirements with the least operational overhead because AWS Amplify is a complete solution that lets frontend web and mobile developers easily build, ship, and host full-stack applications on AWS, with the flexibility to leverage the breadth of AWS services as use cases evolve. No cloud expertise needed1. By using AWS Amplify, the company can refactor the application to a serverless architecture that reduces operational complexity and costs. AWS Amplify offers the following features and benefits:

≫ Amplify Studio: A visual interface that enables you to build and deploy a full-stack app quickly, including frontend UI and backend.

≫ Amplify CLI: A local toolchain that enables you to configure and manage an app backend with just a few commands.

≫ Amplify Libraries: Open-source client libraries that enable you to build cloud-powered mobile and web apps.

≫ Amplify UI Components: Open-source design system with cloud-connected components for building feature-rich apps fast.

≫ Amplify Hosting: Fully managed CI/CD and hosting for fast, secure, and reliable static and server-side rendered apps.

By using AWS Amplify to create a static website for uploads of media files, the company can leverage Amplify Studio to visually build a pixel-perfect UI and connect it to a cloud backend in clicks. By using Amplify Hosting to serve the website through Amazon CloudFront, the company can easily deploy its web app or website to the fast, secure, and reliable AWS content delivery network (CDN), with hundreds of points of presence globally. By using Amazon S3 to store the uploaded media files, the company can benefit from a highly scalable, durable, and cost-effective object storage service that can handle any amount of data2. By using Amazon Cognito to authenticate users, the company can add user sign-up, sign-in, and access control to its web app with a fully managed service that scales to support millions of users3.
The other options are not correct because:

≫ Using AWS Application Migration Service to migrate the application server to Amazon EC2 instances would not refactor the application or accelerate development. AWS Application Migration Service (AWS MGN) is a service that enables you to migrate physical servers, virtual machines (VMs), or cloud servers from any source infrastructure to AWS without requiring agents or specialized tools. However, this would not address the challenges of overutilization and data uploads failures. It would also not reduce operational overhead or costs compared to a serverless architecture.

≫ Creating a static website for uploads of media files and using AWS AppSync to create an API would not be as simple or fast as using AWS Amplify. AWS AppSync is a service that enables you to create flexible APIs for securely accessing, manipulating, and combining data from one or more data sources. However, this would require more configuration and management than using Amplify Studio and Amplify Hosting. It would also not provide authentication features like Amazon Cognito.

≫ Setting up AWS IAM Identity Center (AWS Single Sign-On) to give users the ability to sign in to the application would not be as suitable as using Amazon Cognito. AWS Single Sign-On (AWS SSO) is a service that enables you to centrally manage SSO access and user permissions across multiple AWS accounts and business applications. However, this service is designed for enterprise customers who need to manage access for employees or partners across multiple resources. It is not intended for authenticating end users of web or mobile apps.
References:

≫ https://aws.amazon.com/amplify/

≫ https://aws.amazon.com/s3/

≫ https://aws.amazon.com/cognito/

≫ https://aws.amazon.com/mgn/

≫ https://aws.amazon.com/appsync/

≫ https://aws.amazon.com/single-sign-on/

NEW QUESTION 102
- (Exam Topic 2)
A company has several AWS accounts. A development team is building an automation framework for cloud governance and remediation processes. The automation framework uses AWS Lambda functions in a centralized account. A solutions architect must implement a least privilege permissions policy that allows the Lambda functions to run in each of the company's AWS accounts.
Which combination of steps will meet these requirements? (Choose two.)

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

E. In the centralized account, create an IAM role that has roles of the other accounts as trusted entities.Provide minimal permissions.
F. In the other AWS accounts, create an IAM role that has permissions to assume the role of the centralized accoun
G. Add the Lambda service as a trusted entity.
H. In the other AWS accounts, create an IAM role that has minimal permission
I. Add the Lambda service as a trusted entity.

Answer: AB

Explanation:
https://medium.com/@it.melnichenko/invoke-a-lambda-across-multiple-aws-accounts-8c094b2e70be

NEW QUESTION 104
- (Exam Topic 2)
A company is running a critical stateful web application on two Linux Amazon EC2 instances behind an Application Load Balancer (ALB) with an Amazon RDS for MySQL database The company hosts the DNS records for the application in Amazon Route 53 A solutions architect must recommend a solution to improve the resiliency of the application
The solution must meet the following objectives:
• Application tier RPO of 2 minutes. RTO of 30 minutes
• Database tier RPO of 5 minutes RTO of 30 minutes
The company does not want to make significant changes to the existing application architecture The company must ensure optimal latency after a failover
Which solution will meet these requirements?

A. Configure the EC2 instances to use AWS Elastic Disaster Recovery Create a cross-Region read replica for the RDS DB instance Create an ALB in a second AWS Region Create an AWS Global Accelerator endpoint and associate the endpoint with the ALBs Update DNS records to point to the Global Accelerator endpoint
B. Configure the EC2 instances to use Amazon Data Lifecycle Manager (Amazon DLM) to take snapshots of the EBS volumes Configure RDS automated backups Configure backup replication to a second AWS Region Create an ALB in the second Region Create an AWS Global Accelerator endpoint, and associate the endpoint with the ALBs Update DNS records to point to the Global Accelerator endpoint
C. Create a backup plan in AWS Backup for the EC2 instances and RDS DB instance Configure backup replication to a second AWS Region Create an ALB in the second Region Configure an Amazon CloudFront distribution in front of the ALB Update DNS records to point to CloudFront
D. Configure the EC2 instances to use Amazon Data Lifecycle Manager (Amazon DLM) to take snapshots of the EBS volumes Create a cross-Region read replica for the RDS DB instance Create an ALB in a second AWS Region Create an AWS Global Accelerator endpoint and associate the endpoint with the ALBs

Answer: B

Explanation:
This option meets the RPO and RTO requirements for both the application and database tiers and uses tools like Amazon DLM and RDS automated backups to create and manage the backups. Additionally, it uses Global Accelerator to ensure low latency after failover by directing traffic to the closest healthy endpoint.

NEW QUESTION 107
- (Exam Topic 2)
A software-as-a-service (SaaS) provider exposes APIs through an Application Load Balancer (ALB). The ALB connects to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster that is deployed in the
us-east-I Region. The exposed APIs contain usage of a few non-standard REST methods: LINK, UNLINK, LOCK, and UNLOCK.
Users outside the United States are reporting long and inconsistent response times for these APIs. A solutions architect needs to resolve this problem with a solution that minimizes operational overhead.
Which solution meets these requirements?

A. Add an Amazon CloudFront distributio
B. Configure the ALB as the origin.
C. Add an Amazon API Gateway edge-optimized API endpoint to expose the API
D. Configure the ALB as the target.
E. Add an accelerator in AWS Global Accelerato
F. Configure the ALB as the origin.
G. Deploy the APIs to two additional AWS Regions: eu-west-I and ap-southeast-2. Add latency-based routing records in Amazon Route 53.

Answer: C

Explanation:
Adding an accelerator in AWS Global Accelerator will enable improving the performance of the APIs for local and global users1. AWS Global Accelerator is a service that uses the AWS global network to route traffic to the optimal regional endpoint based on health, client location, and policies1. Configuring the ALB as the origin will enable connecting the accelerator to the ALB that exposes the APIs2. AWS Global Accelerator supports non-standard REST methods such as LINK, UNLINK, LOCK, and UNLOCK3.

NEW QUESTION 108
- (Exam Topic 2)
A company is using AWS Organizations to manage multiple AWS accounts. For security purposes, the company requires the creation of an Amazon Simple Notification Service (Amazon SNS) topic that enables integration with a third-party alerting system in all the Organizations member accounts.
A solutions architect used an AWS CloudFormation template to create the SNS topic and stack sets to automate the deployment of Cloud Formation stacks.
Trusted access has been enabled in Organizations.
What should the solutions architect do to deploy the CloudFormation StackSets in all AWS accounts?

A. Create a stack set in the Organizations member account
B. Use service-managed permission
C. Set deployment options to deploy to an organizatio
D. Use CloudFormation StackSets drift detection.
E. Create stacks in the Organizations member account
F. Use self-service permission
G. Set deploymentoptions to deploy to an organizatio
H. Enable the CloudFormation StackSets automatic deployment.

I. Create a stack set in the Organizations management accoun
J. Use service-managed permission
K. Set deployment options to deploy to the organizatio
L. Enable CloudFormation StackSets automatic deployment.
M. Create stacks in the Organizations management accoun
N. Use service-managed permission
O. Set deployment options to deploy to the organizatio
P. Enable CloudFormation StackSets drift detection.

Answer: C

Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-orgs-manage-auto-deployment.h

NEW QUESTION 109
- (Exam Topic 2)
A company runs an IoT application in the AWS Cloud. The company has millions of sensors that collect data from houses in the United States. The sensors use the MOTT protocol to connect and send data to a custom MQTT broker. The MQTT broker stores the data on a single Amazon EC2 instance. The sensors connect to the broker through the domain named iot.example.com. The company uses Amazon Route 53 as its DNS service. The company stores the data in Amazon DynamoDB.
On several occasions, the amount of data has overloaded the MOTT broker and has resulted in lost sensor data. The company must improve the reliability of the solution.
Which solution will meet these requirements?

A. Create an Application Load Balancer (ALB) and an Auto Scaling group for the MOTT broke
B. Use the Auto Scaling group as the target for the AL
C. Update the DNS record in Route 53 to an alias recor
D. Point the alias record to the AL
E. Use the MQTT broker to store the data.
F. Set up AWS IoT Core to receive the sensor dat
G. Create and configure a custom domain to connect to AWS IoT Cor
H. Update the DNS record in Route 53 to point to the AWS IoT Core Data-ATS endpoin
I. Configure an AWS IoT rule to store the data.
J. Create a Network Load Balancer (NLB). Set the MQTT broker as the targe
K. Create an AWS Global Accelerator accelerato
L. Set the NLB as the endpoint for the accelerato
M. Update the DNS record in Route 53 to a multivalue answer recor
N. Set the Global Accelerator IP addresses as value
O. Use the MQTT broker to store the data.
P. Set up AWS IoT Greengrass to receive the sensor dat
Q. Update the DNS record in Route 53 to point to the AWS IoT Greengrass endpoin
R. Configure an AWS IoT rule to invoke an AWS Lambda function to store the data.

Answer: A

Explanation:
it describes a solution that uses an Application Load Balancer (ALB) and an Auto Scaling group for the MQTT broker. The ALB distributes incoming traffic across the instances in the Auto Scaling group and allows for automatic scaling based on incoming traffic. The use of an alias record in Route 53 allows for easy updates to the DNS record without changing the IP address. This solution improves the reliability of the MQTT broker by allowing it to automatically scale based on incoming traffic, reducing the likelihood of lost data due to broker overload.
Reference: https://aws.amazon.com/elasticloadbalancing/applicationloadbalancer/ https://aws.amazon.com/autoscaling/ https://aws.amazon.com/route53/

NEW QUESTION 113
- (Exam Topic 2)
A company needs to build a disaster recovery (DR) solution for its ecommerce website. The web application is hosted on a fleet of t3.large Amazon EC2 instances and uses an Amazon RDS for MySQL DB instance. The EC2 instances are in an Auto Scaling group that extends across multiple Availability Zones.
In the event of a disaster, the web application must fail over to the secondary environment with an RPO of 30 seconds and an R TO of 10 minutes.
Which solution will meet these requirements MOST cost-effectively?

A. Use infrastructure as code (IaC) to provision the new infrastructure in the DR Regio
B. Create a cross-Region read replica for the DB instanc
C. Set up a backup plan in AWS Backup to createcross-Region backups for the EC2 instances and the DB instanc
D. Create a cron expression to back up the EC2 instances and the DB instance every 30 seconds to the DR Regio
E. Recover the EC2 instancesfrom the latest EC2 backu
F. Use an Amazon Route 53 geolocation routing policy to automatically fail over to the DR Region in the event of a disaster.
G. Use infrastructure as code (IaC) to provision the new infrastructure in the DR Regio
H. Create across-Region read replica for the DB instanc
I. Set up AWS Elastic Disaster Recovery to continuously replicate the EC2 instances to the DR Regio
J. Run the EC2 instances at the minimum capacity in the DR Region Use an Amazon Route 53 failover routing policy to automatically fail over to the DR Region in the event of a disaste
K. Increase the desired capacity of the Auto Scaling group.
L. Set up a backup plan in AWS Backup to create cross-Region backups for the EC2 instances and the DB instanc
M. Create a cron expression to back up the EC2 instances and the DB instance every 30 seconds to the DR Regio
N. Use infrastructure as code (IaC) to provision the new infrastructure in the DR Regio
O. Manually restore the backed-up data on new instance
P. Use an Amazon Route 53 simple routing policy to automatically fail over to the DR Region in the event of a disaster.
Q. Use infrastructure as code (IaC) to provision the new infrastructure in the DR Regio
R. Create an Amazon Aurora global databas
S. Set up AWS Elastic Disaster Recovery to continuously replicate the EC2 instances to the DR Regio
T. Run the Auto Scaling group of EC2 instances at full capacity in the DR Regio
. Use an Amazon Route 53 failover routing policy to automatically fail over to the DR Region in the event of a disaster.

Answer: B

Explanation:
The company should use infrastructure as code (IaC) to provision the new infrastructure in the DR Region. The company should create a cross-Region read replica for the DB instance. The company should set up AWS Elastic Disaster Recovery to continuously replicate the EC2 instances to the DR Region. The company should run the EC2 instances at the minimum capacity in the DR Region. The company should use an Amazon Route 53 failover routing policy to automatically fail over to the DR Region in the event of a disaster. The company should increase the desired capacity of the Auto Scaling group. This solution will meet the requirements most cost-effectively because AWS Elastic Disaster Recovery (AWS DRS) is a service that minimizes downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications using affordable storage, minimal compute, and point-in-time recovery. AWS DRS enables RPOs of seconds and RTOs of minute1s. AWS DRS continuously replicates data from the source servers to a staging area subnet in the DR Region, where it uses low-cost storage and minimal compute resources to maintain ongoing replication. In the event of a disaster, AWS DRS automatically converts the servers to boot and run natively on AWS and launches recovery instances on AWS within minutes2. By using AWS DRS, the company can save costs by removing idle recovery site resources and paying for the full disaster recovery site only when needed. By creating a cross-Region read replica for the DB instance, the company can have a standby copy of its primary database in a different AWS Region3. By using infrastructure as code (IaC), the company can provision the new infrastructure in the DR Region in an automated and consistent way4. By using an Amazon Route 53 failover routing policy, the company can route traffic to a resource that is healthy or to another resource when the first resource becomes unavailable.
The other options are not correct because:

≫ Using AWS Backup to create cross-Region backups for the EC2 instances and the DB instance would not meet the RPO and RTO requirements. AWS Backup is a service that enables you to centralize and automate data protection across AWS services. You can use AWS Backup to back up your application data across AWS services in your account and across accounts. However, AWS Backup does not
provide continuous replication or fast recovery; it creates backups at scheduled intervals and requires manual restoration. Creating backups every 30 seconds would also incur high costs and network bandwidth.

≫ Creating an Amazon API Gateway Data API service integration with Amazon Redshift would not help with disaster recovery. The Data API is a feature that enables you to query your Amazon Redshift cluster using HTTP requests, without needing a persistent connection or a SQL client. It is useful for building applications that interact with Amazon Redshift, but not for replicating or recovering data.

≫ Creating an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster would not help with disaster recovery. AWS Data Exchange is a service that makes it easy for AWS customers to exchange data in the cloud. You can use AWS Data Exchange to subscribe to a diverse selection of third-party data products or offer your own data products to other AWS customers. A datashare is a feature that enables you to share live and secure access to your Amazon Redshift data across your accounts or with third parties without copying or moving the underlying data. It is useful for sharing query results and views with other users, but not for replicating or recovering data.
References:

≫ https://aws.amazon.com/disaster-recovery/

≫ https://docs.aws.amazon.com/drs/latest/userguide/what-is-drs.html

≫ https://aws.amazon.com/cloudformation/

≫ https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover.html

≫ https://aws.amazon.com/backup/

≫ https://docs.aws.amazon.com/redshift/latest/mgmt/data-api.html

≫ https://aws.amazon.com/data-exchange/

≫ https://docs.aws.amazon.com/redshift/latest/dg/datashare-overview.html

NEW QUESTION 118
- (Exam Topic 2)
A company wants to refactor its retail ordering web application that currently has a load-balanced Amazon EC2 instance fleet for web hosting, database API services, and business logic. The company needs to create a decoupled, scalable architecture with a mechanism for retaining failed orders while also minimizing operational costs.
Which solution will meet these requirements?

A. Use Amazon S3 for web hosting with Amazon API Gateway for database API service
B. Use Amazon Simple Queue Service (Amazon SQS) for order queuin
C. Use Amazon Elastic Container Service (Amazon ECS) for business logic with Amazon SQS long polling for retaining failed orders.
D. Use AWS Elastic Beanstalk for web hosting with Amazon API Gateway for database API service
E. Use Amazon MQ for order queuin
F. Use AWS Step Functionsfor business logic with Amazon S3 Glacier Deep Archive for retaining failed orders.
G. Use Amazon S3 for web hosting with AWS AppSync for database API service
H. Use Amazon Simple Queue Service (Amazon SQS) for order queuin
I. Use AWS Lambda for business logic with an Amazon SQS dead-letter queue for retaining failed orders.
J. Use Amazon Lightsail for web hosting with AWS AppSync for database API service
K. Use Amazon Simple Email Service (Amazon SES) for order queuin
L. UseAmazon Elastic Kubernetes Service (Amazon EKS) for business logic with Amazon OpenSearch Service for retaining failed orders.

Answer: C

Explanation:
•Use Amazon S3 for web hosting with AWS AppSync for database API services. Use Amazon Simple Queue Service (Amazon SQS) for order queuing. Use AWS Lambda for business logic with an Amazon SQS dead-letter queue for retaining failed orders.
This solution will allow you to:
•Host a static website on Amazon S3 without provisioning or managing servers1.
•Use AWS AppSync to create a scalable GraphQL API that connects to your database and other data sources1.
•Use Amazon SQS to decouple and scale your order processing microservices1.
•Use AWS Lambda to run code for your business logic without provisioning or managing servers1.
•Use an Amazon SQS dead-letter queue to retain messages that can't be processed by your Lambda function1.

NEW QUESTION 123
- (Exam Topic 2)
A company uses a load balancer to distribute traffic to Amazon EC2 instances in a single Availability Zone. The company is concerned about security and wants a solutions architect to re-architect the solution to meet the following requirements:
• Inbound requests must be filtered for common vulnerability attacks.

• Rejected requests must be sent to a third-party auditing application.
• All resources should be highly available. Which solution meets these requirements?

A. Configure a Multi-AZ Auto Scaling group using the application's AM
B. Create an Application Load Balancer (ALB) and select the previously created Auto Scaling group as the targe
C. Use Amazon Inspector to monitor traffic to the ALB and EC2 instance
D. Create a web ACL in WA
E. Create an AWS WAF using the web ACL and AL
F. Use an AWS Lambda function to frequently push the Amazon Inspector report to the third-party auditing application.
G. Configure an Application Load Balancer (ALB) and add the EC2 instances as targets Create a web ACL in WA
H. Create an AWS WAF using the web ACL and ALB name and enable logging with Amazon CloudWatch Log
I. Use an AWS Lambda function to frequently push the logs to the third-party auditing application.
J. Configure an Application Load Balancer (ALB) along with a target group adding the EC2 instances as target
K. Create an Amazon Kinesis Data Firehose with the destination of the third-party auditing applicatio
L. Create a web ACL in WA
M. Create an AWS WAF using the web ACL and ALB then enable logging by selecting the Kinesis Data Firehose as the destinatio
N. Subscribe to AWS Managed Rules in AWS Marketplace, choosing the WAF as the subscriber.
O. Configure a Multi-AZ Auto Scaling group using the application's AM
P. Create an Application Load Balancer (ALB) and select the previously created Auto Scaling group as the targe
Q. Create an Amazon Kinesis Data Firehose with a destination of the third-party auditing applicatio
R. Create a web ACL inWA
S. Create an AWS WAF using the WebACL and ALB then enable logging by selecting the Kinesis Data Firehose as the destinatio
T. Subscribe to AWS Managed Rules in AWS Marketplace, choosing the WAF as the subscriber.

Answer: D

Explanation:
https://docs.aws.amazon.com/waf/latest/developerguide/marketplace-managed-rule-groups.html

NEW QUESTION 127
- (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. Enable Aurora Auto Scaling for Aurora Replica
F. Use an Application Load Balancer with the round robin routing and sticky sessions enabled.
G. Enable Aurora Scaling for Aurora writer
H. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.

Answer: C

Explanation:
Aurora Auto Scaling enables your Aurora DB cluster to handle sudden increases in connectivity or workload. When the connectivity or workload decreases, Aurora Auto Scaling removes unnecessary Aurora Replicas so that you don't pay for unused provisioned DB instances

NEW QUESTION 129
- (Exam Topic 2)
A company wants to migrate to AWS. The company is running thousands of VMs in a VMware ESXi environment. The company has no configuration management database and has little Knowledge about the utilization of the VMware portfolio.
A solutions architect must provide the company with an accurate inventory so that the company can plan for a cost-effective migration.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use AWS Systems Manager Patch Manager to deploy Migration Evaluator to each V
B. Review the collected data in Amazon QuickSigh
C. Identify servers that have high utilizatio
D. Remove the servers that have high utilization from the migration lis
E. Import the data to AWS Migration Hub.
F. Export the VMware portfolio to a csv fil
G. Check the disk utilization for each serve
H. Remove servers that have high utilizatio
I. Export the data to AWS Application Migration Servic
J. Use AWS Server Migration Service (AWS SMS) to migrate the remaining servers.
K. Deploy the Migration Evaluator agentless collector to the ESXi hyperviso
L. Review the collected data in Migration Evaluato
M. Identify inactive server
N. Remove the inactive servers from the migration lis
O. Import the data to AWS Migration Hub.
P. Deploy the AWS Application Migration Service Agent to each V
Q. When the data is collected, useAmazon Redshift to import and analyze the dat
R. Use Amazon QuickSight for data visualization.

Answer: C

Explanation:

https://aws.amazon.com/migration-evaluator/features/

NEW QUESTION 131
- (Exam Topic 2)
A manufacturing company is building an inspection solution for its factory. The company has IP cameras at the end of each assembly line. The company has used Amazon SageMaker to train a machine learning (ML) model to identify common defects from still images.

The company wants to provide local feedback to factory workers when a defect is detected. The company must be able to provide this feedback even if the factory's internet connectivity is down. The company has a local Linux server that hosts an API that provides local feedback to the workers.

How should the company deploy the ML model to meet these requirements?

A. Set up an Amazon Kinesis video stream from each IP camera to AW
B. Use Amazon EC2 instances to take still images of the stream
C. Upload the images to an Amazon S3 bucke
D. Deploy a SageMaker endpoint with the ML mode
E. Invoke an AWS Lambda function to call the inference endpoint when new images are uploade
F. Configure the Lambda function to call the local API when a defect is detected.
G. Deploy AWS IoT Greengrass on the local serve
H. Deploy the ML model to the Greengrass serve
I. Create a Greengrass component to take still images from the cameras and run inferenc
J. Configure the component to call the local API when a defect is detected.
K. Order an AWS Snowball devic
L. Deploy a SageMaker endpoint the ML model and an Amazon EC2 instance on the Snowball devic
M. Take still images from the camera
N. Run inference from the EC2 instanc
O. Configure the instance to call the local API when a defect is detected.
P. Deploy Amazon Monitron devices on each IP camer
Q. Deploy an Amazon Monitron Gateway on premise
R. Deploy the ML model to the Amazon Monitron device
S. Use Amazon Monitron health state alarms to call the local API from an AWS Lambda function when a defect is detected.

Answer: B

Explanation:
The company should use AWS IoT Greengrass to deploy the ML model to the local server and provide local feedback to the factory workers. AWS IoT Greengrass is a service that extends AWS cloud capabilities to local devices, allowing them to collect and analyze data closer to the source of information, react autonomously to local events, and communicate securely with each other on local networks1. AWS IoT
Greengrass also supports ML inference at the edge, enabling devices to run ML models locally without requiring internet connectivity2.
The other options are not correct because:

> Setting up an Amazon Kinesis video stream from each IP camera to AWS would not work if the factory's internet connectivity is down. It would also incur unnecessary costs and latency to stream video data to the cloud and back.

> Ordering an AWS Snowball device would not be a scalable or cost-effective solution for deploying the ML model. AWS Snowball is a service that provides physical devices for data transfer and edge computing, but it is not designed for continuous operation or frequent updates3.

> Deploying Amazon Monitron devices on each IP camera would not work because Amazon Monitron is a service that monitors the condition and performance of industrial equipment using sensors and machine learning, not cameras4.
References:

> https://aws.amazon.com/greengrass/

> https://docs.aws.amazon.com/greengrass/v2/developerguide/use-machine-learning-inference.html

> https://aws.amazon.com/snowball/

> https://aws.amazon.com/monitron/

NEW QUESTION 135
- (Exam Topic 2)
A company is running an application in the AWS Cloud. The core business logic is running on a set of Amazon EC2 instances in an Auto Scaling group. An Application Load Balancer (ALB) distributes traffic to the EC2 instances. Amazon Route 53 record api.example.com is pointing to the ALB.

The company's development team makes major updates to the business logic. The company has a rule that when changes are deployed, only 10% of customers can receive the new logic during a testing window. A customer must use the same version of the business logic during the testing window.

How should the company deploy the updates to meet these requirements?

A. Create a second ALB, and deploy the new logic to a set of EC2 instances in a new Auto Scaling group.Configure the ALB to distribute traffic to the EC2 instance
B. Update the Route 53 record to use weighted routing, and point the record to both of the ALBs.
C. Create a second target group that is referenced by the AL
D. Deploy the new logic to EC2 instances in this new target grou
E. Update the ALB listener rule to use weighted target group
F. Configure ALB target group stickiness.
G. Create a new launch configuration for the Auto Scaling grou
H. Specify the launch configuration to use the AutoScalingRoIIingUpdate policy, and set the MaxBatchSize option to 10. Replace the launch configuration on the Auto Scaling grou
I. Deploy the changes.
J. Create a second Auto Scaling group that is referenced by the AL
K. Deploy the new logic on a set ofEC2 instances in this new Auto Scaling grou
L. Change the ALB routing algorithm to least outstanding requests (LOR). Configure ALB session stickiness.

Answer: B

Explanation:
The company should create a second target group that is referenced by the ALB. The company should deploy the new logic to EC2 instances in this new target group. The company should update the ALB listener rule to use weighted target groups. The company should configure ALB target group stickiness. This solution will meet the requirements because weighted target groups are a feature that enables you to distribute traffic across multiple target groups using a single listener

rule. You can specify a weight for each target group, which determines the percentage of requests that are routed to that target group. For example, if you specify two target groups, each with a weight of 10, each target group receives half the requests1. By creating a second target group and deploying the new logic to EC2 instances in this new target group, the company can have two versions of its business logic running in parallel. By updating the ALB listener rule to use weighted target groups, the company can control how much traffic is sent to each version. By configuring ALB target group stickiness, the company can ensure that a customer uses the same version of the business logic during the testing window. Target group stickiness is a feature that enables you to bind a user's session to a specific target within a target group for the duration of the session2.

The other options are not correct because:

≫ Creating a second ALB and deploying the new logic to a set of EC2 instances in a new Auto Scaling group would not be as cost-effective or simple as using weighted target groups. A second ALB would incur additional charges and require more configuration and management. Updating the Route 53 record to use weighted routing would not ensure that a customer uses the same version of the business logic during the testing window, as DNS caching could affect how requests are routed.

≫ Creating a new launch configuration for the Auto Scaling group and replacing it on the Auto Scaling group would not allow for gradual traffic shifting between versions. A launch configuration is a
template that an Auto Scaling group uses to launch EC2 instances. You can specify information such as
the AMI ID, instance type, key pair, security groups, and block device mapping for your instances3. However, replacing the launch configuration on an Auto Scaling group would affect all instances in that group, not just 10% of customers.

≫ Creating a second Auto Scaling group and changing the ALB routing algorithm to least outstanding requests (LOR) would not allow for controlled traffic shifting between versions. A second Auto Scaling group would require more configuration and management. The LOR routing algorithm is a feature that enables you to route traffic based on how quickly targets respond to requests. The load balancer selects a target from the target group with the fewest outstanding requests4. However, this algorithm does not
take into account customer sessions or weights.

References:

≫ https://docs.aws.amazon.com/elasticloadbalancing/latest/application/sticky-sessions.html

≫ https://docs.aws.amazon.com/autoscaling/ec2/userguide/LaunchConfiguration.html

≫ https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-target-groups.html#rou

NEW QUESTION 138
- (Exam Topic 2)
A company is using an organization in AWS Organizations to manage hundreds of AWS accounts. A solutions architect is working on a solution to provide baseline protection for the Open Web Application Security Project (OWASP) top 10 web application vulnerabilities. The solutions architect is using AWS WAF for all existing and new Amazon CloudFront distributions that are deployed within the organization.
Which combination of steps should the solutions architect take to provide the baseline protection? (Select THREE.)

A. Enable AWS Config in all accounts.
B. Enable Amazon GuardDuty in all accounts.
C. Enable all features for the organization.
D. Use AWS Firewall Manager to deploy AWS WAF rules in all accounts for all CloudFront distributions.
E. Use AWS Shield Advanced to deploy AWS WAF rules in all accounts for all CloudFront distributions.
F. Use AWS Security Hub to deploy AWS WAF rules in all accounts for all CloudFront distributions.

Answer: CDE

Explanation:
Enabling all features for the organization will enable using AWS Firewall Manager to centrally configure and manage firewall rules across multiple AWS accounts1. Using AWS Firewall Manager to deploy AWS WAF rules in all accounts for all CloudFront distributions will enable providing baseline protection for the OWASP top 10 web application vulnerabilities2. AWS Firewall Manager supports AWS WAF rules that can help protect against common web exploits such as SQL injection and cross-site scripting3. Configuring redirection of HTTP requests to HTTPS requests in CloudFront will enable encrypting the data in transit using SSL/TLS.

NEW QUESTION 141
- (Exam Topic 3)
A company wants to migrate its on-premises application to AWS. The database for the application stores structured product data and temporary user session data. The company needs to decouple the product data from the user session data. The company also needs to implement replication in another AWS Region for disaster recovery.
Which solution will meet these requirements with the HIGHEST performance?

A. Create an Amazon RDS DB instance with separate schemas to host the product data and the user session dat
B. Configure a read replica for the DB instance in another Region.
C. Create an Amazon RDS DB instance to host the product dat
D. Configure a read replica for the DB instance in another Regio
E. Create a global datastore in Amazon ElastiCache for Memcached to host the user session data.
F. Create two Amazon DynamoDB global table
G. Use one global table to host the product data Use the other global table to host the user session dat
H. Use DynamoDB Accelerator (DAX) for caching.
I. Create an Amazon RDS DB instance to host the product dat
J. Configure a read replica for the DB instance in another Regio
K. Create an Amazon DynamoDB global table to host the user session data

Answer: B

NEW QUESTION 142
- (Exam Topic 3)
A company needs to create and manage multiple AWS accounts for a number of departments from a central location. The security team requires read-only access to all accounts from its own AWS account. The company is using AWS Organizations and created an account for the security team.
How should a solutions architect meet these requirements?

A. Use the OrganizationAccountAccessRole IAM role to create a new IAM policy with read-only access in each member accoun
B. Establish a trust relationship between the IAM policy in each member account and the security accoun
C. Ask the security team to use the IAM policy to gain access.

D. Use the Organization AccountAccessRole IAM role to create a new IAM role with read-only access in each member accoun
E. Establish a trust relationship between the IAM role in each member account and the security accoun
F. Ask the security team to use the IAM role to gain access.
G. Ask the security team to use AWS Security Token Service (AWS STS) lo call the AssumeRole API tor the Organization AccountAccessRole IAM role in the management account from the security accoun
H. Use the generated temporary credentials to gain access.
I. Ask the security team to use AWS Security Token Service (AWS STS) to call the AssumeRole API for the Organization AccountAccessRole IAM role in the member account from the security accoun
J. Use the generated temporary credentials to gain access.

Answer: B

Explanation:
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_access.html#orgs_manage_ "When you create a member account using the AWS Organizations console, AWS Organizations
automatically creates an IAM role named OrganizationAccountAccessRole in the account" you need
OrganizationAccountAccessRole in member account to create an read-only role and use role from security team to assume this read-only role.

NEW QUESTION 146
- (Exam Topic 3)
A solutions architect must update an application environment within AWS Elastic Beanstalk using a blue/green deployment methodology The solutions architect creates an environment that is identical to the existing application environment and deploys the application to the new environment.
What should be done next to complete the update?

A. Redirect to the new environment using Amazon Route 53
B. Select the Swap Environment URLs option
C. Replace the Auto Scaling launch configuration
D. Update the DNS records to point to the green environment

Answer: B

Explanation:
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.CNAMESwap.html

NEW QUESTION 148
- (Exam Topic 3)
A company has developed a mobile game. The backend for the game runs on several virtual machines located in an on-premises data center. The business logic is exposed using a REST API with multiple functions. Player session data is stored in central file storage. Backend services use different API keys for throttling and to distinguish between live and test traffic.
The load on the game backend varies throughout the day. During peak hours, the server capacity is not sufficient. There are also latency issues when fetching player session data. Management has asked a solutions architect to present a cloud architecture that can handle the game's varying load and provide low-latency data access. The API model should not be changed.
Which solution meets these requirements?

A. Implement the REST API using a Network Load Balancer (NLB). Run the business logic on an Amazon EC2 instance behind the NL
B. Store player session data in Amazon Aurora Serverless.
C. Implement the REST API using an Application Load Balancer (ALB). Run the business logic in AWS Lambd
D. Store player session data in Amazon DynamoDB with on-demand capacity.
E. Implement the REST API using Amazon API Gatewa
F. Run the business logic in AWS Lambd
G. Store player session data in Amazon DynamoDB with on- demand capacity.
H. Implement the REST API using AWS AppSyn
I. Run the business logic in AWS Lambd
J. Store player session data in Amazon Aurora Serverless.

Answer: C

NEW QUESTION 149
- (Exam Topic 3)
A company wants to migrate its website from an on-premises data center onto AWS. At the same time, it wants to migrate the website to a containerized microservice-based architecture to improve the availability and cost efficiency. The company's security policy states that privileges and network permissions must be configured according to best practice, using least privilege.
A Solutions Architect must create a containerized architecture that meets the security requirements and has deployed the application to an Amazon ECS cluster.
What steps are required after the deployment to meet the requirements? (Choose two.)

A. Create tasks using the bridge network mode.
B. Create tasks using the awsvpc network mode.
C. Apply security groups to Amazon EC2 instances, and use IAM roles for EC2 instances to access other resources.
D. Apply security groups to the tasks, and pass IAM credentials into the container at launch time to access other resources.
E. Apply security groups to the tasks, and use IAM roles for tasks to access other resources.

Answer: BE

Explanation:
The awsvpc network mode provides each task with its own elastic network interface (ENI) and a primary private IP address1. By using this network mode, the solutions architect can isolate the tasks from each other and apply security groups to the tasks directly2. This way, the solutions architect can control the inbound and outbound traffic at the task level and enforce the least privilege principle3. IAM roles for tasks allow the solutions architect to assign permissions to each task separately, so that they can access other AWS resources that they need4. By using IAM roles for tasks, the solutions architect can avoid passing IAM credentials into the container at launch time, which is less secure and more prone to errors5.
References:

- awsvpc network mode
- Task networking with the awsvpc network mode
- Security groups for your VPC
- IAM roles for tasks
- Best practices for managing AWS access keys

NEW QUESTION 154
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual SAP-C02 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the SAP-C02 Product From:

https://www.2passeasy.com/dumps/SAP-C02/

Money Back Guarantee

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