SAP-C02 Dumps

AWS Certified Solutions Architect - Professional

https://www.certleader.com/SAP-C02-dumps.html

NEW QUESTION 1
- (Exam Topic 1)
A company is developing a new serverless API by using Amazon API Gateway and AWS Lambda. The company integrated the Lambda functions with API Gateway to use several shared libraries and custom classes.
A solutions architect needs to simplify the deployment of the solution and optimize for code reuse. Which solution will meet these requirements?

A. Deploy the shared libraries and custom classes into a Docker imag
B. Store the image in an S3 bucket.Create a Lambda layer that uses the Docker image as the sourc
C. Deploy the API's Lambda functions as Zip package
D. Configure the packages to use the Lambda layer.
E. Deploy the shared libraries and custom classes to a Docker imag
F. Upload the image to Amazon Elastic Container Registry (Amazon ECR). Create a Lambda layer that uses the Docker image as the sourc
G. Deploy the API's Lambda functions as Zip package
H. Configure the packages to use the Lambda layer.
I. Deploy the shared libraries and custom classes to a Docker container in Amazon Elastic Container Service (Amazon ECS) by using the AWS Fargate launch typ
J. Deploy the API's Lambda functions as Zip package
K. Configure the packages to use the deployed container as a Lambda layer.
L. Deploy the shared libraries, custom classes, and code for the API's Lambda functions to a Docker imag
M. Upload the image to Amazon Elastic Container Registry (Amazon ECR). Configure the API's Lambda functions to use the Docker image as the deployment package.

Answer: B

Explanation:
Deploying the shared libraries and custom classes to a Docker image and uploading the image to Amazon Elastic Container Registry (Amazon ECR) and creating a Lambda layer that uses the Docker image as the source. Then, deploying the API's Lambda functions as Zip packages and configuring the packages to use the Lambda layer would meet the requirements for simplifying the deployment and optimizing for code reuse.
A Lambda layer is a distribution mechanism for libraries, custom runtimes, and other function dependencies. It allows you to manage your in-development function code separately from your dependencies, this way you can easily update your dependencies without having to update your entire function code.
By deploying the shared libraries and custom classes to a Docker image and uploading the image to Amazon Elastic Container Registry (ECR), it makes it easy to manage and version the dependencies. This way, the company can use the same version of the dependencies across different Lambda functions.
By creating a Lambda layer that uses the Docker image as the source, the company can configure the API's Lambda functions to use the layer, reducing the need to include the dependencies in each function package, and making it easy to update the dependencies across all functions at once.
Reference:
AWS Lambda Layers documentation: https://docs.aws.amazon.com/lambda/latest/dg/configuration-layers.html
AWS Elastic Container Registry (ECR) documentation: https://aws.amazon.com/ecr/ Building Lambda Layers with Docker documentation:
https://aws.amazon.com/blogs/compute/building-lambda-layers-with-docker/

NEW QUESTION 2
- (Exam Topic 1)
An application is using an Amazon RDS for MySQL Multi-AZ DB instance in the us-east-1 Region. After a failover test, the application lost the connections to the database and could not re-establish the connections. After a restart of the application, the application re-established the connections.
A solutions architect must implement a solution so that the application can re-establish connections to the database without requiring a restart.
Which solution will meet these requirements?

A. Create an Amazon Aurora MySQL Serverless v1 DB instanc
B. Migrate the RDS DB instance to the Aurora Serverless v1 DB instanc
C. Update the connection settings in the application to point to the Aurora reader endpoint.
D. Create an RDS prox
E. Configure the existing RDS endpoint as a targe
F. Update the connection settings in the application to point to the RDS proxy endpoint.
G. Create a two-node Amazon Aurora MySQL DB cluste
H. Migrate the RDS DB instance to the Aurora DB cluste
I. Create an RDS prox
J. Configure the existing RDS endpoint as a targe
K. Update the connection settings in the application to point to the RDS proxy endpoint.
L. Create an Amazon S3 bucke
M. Export the database to Amazon S3 by using AWS Database Migration Service (AWS DMS). Configure Amazon Athena to use the S3 bucket as a data stor
N. Install the latest Open Database Connectivity (ODBC) driver for the applicatio
O. Update the connection settings in the application to point to the Athena endpoint

Answer: B

Explanation:
Amazon RDS Proxy is a fully managed database proxy service for Amazon Relational Database Service (RDS) that makes applications more scalable, resilient, and secure. It allows applications to pool and share connections to an RDS database, which can help reduce database connection overhead, improve scalability, and provide automatic failover and high availability.

NEW QUESTION 3
- (Exam Topic 1)
A large mobile gaming company has successfully migrated all of its on-premises infrastructure to the AWS Cloud. A solutions architect is reviewing the environment to ensure that it was built according to the design and that it is running in alignment with the Well-Architected Framework.
While reviewing previous monthly costs in Cost Explorer, the solutions architect notices that the creation and subsequent termination of several large instance types account for a high proportion of the costs. The solutions architect finds out that the company's developers are launching new Amazon EC2 instances as part of their testing and that the developers are not using the appropriate instance types.
The solutions architect must implement a control mechanism to limit the instance types that only the developers can launch.
Which solution will meet these requirements?

A. Create a desired-instance-type managed rule in AWS Confi

B. Configure the rule with the instance types that are allowe
C. Attach the rule to an event to run each time a new EC2 instance is launched.
D. In the EC2 console, create a launch template that specifies the instance types that are allowe
E. Assign the launch template to the developers' IAM accounts.
F. Create a new IAM polic
G. Specify the instance types that are allowe
H. Attach the policy to an IAM group that contains the IAM accounts for the developers
I. Use EC2 Image Builder to create an image pipeline for the developers and assist them in the creation of a golden image.

Answer: C

Explanation:
This is doable with IAM policy creation to restrict users to specific instance types. Found the below article. https://blog.vizuri.com/limiting-allowed-aws-instance-type-with-iam-policy

NEW QUESTION 4
- (Exam Topic 1)
A company is refactoring its on-premises order-processing platform in the AWS Cloud. The platform includes a web front end that is hosted on a fleet of VMs RabbitMQ to connect the front end to the backend, and a Kubernetes cluster to run a containerized backend system to process the orders. The company does not want to make any major changes to the application
Which solution will meet these requirements with the LEAST operational overhead?

A. Create an AMI of the web server VM Create an Amazon EC2 Auto Scaling group that uses the AMI and an Application Load Balancer Set up Amazon MQ to replace the on-premises messaging queue Configure Amazon Elastic Kubernetes Service (Amazon EKS) to host the order-processing backend
B. Create a custom AWS Lambda runtime to mimic the web server environment Create an Amazon API Gateway API to replace the front-end web servers Set up Amazon MQ to replace the on-premises messaging queue Configure Amazon Elastic Kubernetes Service (Amazon EKS) to host theorder-processing backend
C. Create an AMI of the web server VM Create an Amazon EC2 Auto Scaling group that uses the AMI and an Application Load Balancer Set up Amazon MQ to replace the on-premises messaging queue Install Kubernetes on a fleet of different EC2 instances to host the order-processing backend
D. Create an AMI of the web server VM Create an Amazon EC2 Auto Scaling group that uses the AMI and an Application Load Balancer Set up an Amazon Simple Queue Service (Amazon SQS) queue to replace the on-premises messaging queue Configure Amazon Elastic Kubernetes Service (Amazon EKS) to host the order-processing backend

Answer: A

Explanation:
https://aws.amazon.com/about-aws/whats-new/2020/11/announcing-amazon-mq-rabbitmq/

NEW QUESTION 5
- (Exam Topic 1)
An international delivery company hosts a delivery management system on AWS. Drivers use the system to upload confirmation of delivery. Confirmation includes the recipient's signature or a photo of the package with the recipient. The driver's handheld device uploads signatures and photos through FTP to a single Amazon EC2 instance. Each handheld device saves a file in a directory based on the signed-in user, and the file name matches the delivery number. The EC2 instance then adds metadata to the file after querying a central database to pull delivery information. The file is then placed in Amazon S3 for archiving.
As the company expands, drivers report that the system is rejecting connections. The FTP server is having problems because of dropped connections and memory issues. In response to these problems, a system engineer schedules a cron task to reboot the EC2 instance every 30 minutes. The billing team reports that files are not always in the archive and that the central system is not always updated.
A solutions architect needs to design a solution that maximizes scalability to ensure that the archive always receives the files and that systems are always updated. The handheld devices cannot be modified, so the company cannot deploy a new application.
Which solution will meet these requirements?

A. Create an AMI of the existing EC2 instanc
B. Create an Auto Scaling group of EC2 instances behind an Application Load Balance
C. Configure the Auto Scaling group to have a minimum of three instances.
D. Use AWS Transfer Family to create an FTP server that places the files in Amazon Elastic File System (Amazon EFS). Mount the EFS volume to the existing EC2 instanc
E. Point the EC2 instance to the new path for file processing.
F. Use AWS Transfer Family to create an FTP server that places the files in Amazon S3. Use an S3 event notification through Amazon Simple Notification Service (Amazon SNS) to invoke an AWS Lambda functio
G. Configure the Lambda function to add the metadata and update the delivery system.
H. Update the handheld devices to place the files directly in Amazon S3. Use an S3 event notification through Amazon Simple Queue Service (Amazon SQS) to invoke an AWS Lambda functio
I. Configure the Lambda function to add the metadata and update the delivery system.

Answer: C

Explanation:
Using AWS Transfer Family to create an FTP server that places the files in Amazon S3 and using S3 event notifications through Amazon Simple Notification Service (Amazon SNS) to invoke an AWS Lambda function will ensure that the archive always receives the files and that the central system is always updated. This solution maximizes scalability and eliminates the need for manual intervention, such as rebooting the EC2 instance.

NEW QUESTION 6
- (Exam Topic 1)
A company is developing and hosting several projects in the AWS Cloud. The projects are developed across multiple AWS accounts under the same organization in AWS Organizations. The company requires the cost lor cloud infrastructure to be allocated to the owning project. The team responsible for all of the AWS accounts has discovered that several Amazon EC2 instances are lacking the Project tag used for cost allocation.
Which actions should a solutions architect take to resolve the problem and prevent it from happening in the future? (Select THREE.)

A. Create an AWS Config rule in each account to find resources with missing tags.
B. Create an SCP in the organization with a deny action for ec2:RunInstances if the Project tag is missing.
C. Use Amazon Inspector in the organization to find resources with missing tags.

D. Create an IAM policy in each account with a deny action for ec2:RunInstances if the Project tag is missing.
E. Create an AWS Config aggregator for the organization to collect a list of EC2 instances with the missing Project tag.
F. Use AWS Security Hub to aggregate a list of EC2 instances with the missing Project tag.

Answer: ABE

Explanation:
https://docs.aws.amazon.com/config/latest/developerguide/config-rule-multi-account-deployment.html
https://docs.aws.amazon.com/config/latest/developerguide/aggregate-data.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_tagging.htm

NEW QUESTION 7
- (Exam Topic 1)
A company wants to migrate its workloads from on premises to AWS. The workloads run on Linux and Windows. The company has a large on-premises intra structure that consists of physical machines and VMs that host numerous applications.
The company must capture details about the system configuration. system performance. running processure and network coi.net lions of its o. -premises ,on boards. The company also must divide the on-premises applications into groups for AWS migrations. The company needs recommendations for Amazon EC2 instance types so that the company can run its workloads on AWS in the most cost-effective manner.
Which combination of steps should a solutions architect take to meet these requirements? (Select THREE.)

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

Answer: ADE

Explanation:
https://docs.aws.amazon.com/application-discovery/latest/userguide/discovery-agent.html
https://docs.aws.amazon.com/migrationhub/latest/ug/ec2-recommendations.html

NEW QUESTION 8
- (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 9
- (Exam Topic 1)
A company is building a solution in the AWS Cloud. Thousands or devices will connect to the solution and send data. Each device needs to be able to send and receive data in real time over the MQTT protocol. Each device must authenticate by using a unique X.509 certificate.
Which solution will meet these requirements with the LEAST operational overhead?

A. Set up AWS loT Cor
B. For each device, create a corresponding Amazon MQ queue and provision a certificat
C. Connect each device to Amazon MQ.
D. Create a Network Load Balancer (NLB) and configure it with an AWS Lambda authorize
E. Run an MQTT broker on Amazon EC2 instances in an Auto Scaling grou
F. Set the Auto Scaling group as the target for the NL
G. Connect each device to the NLB.
H. Set up AWS loT Cor
I. For each device, create a corresponding AWS loT thing and provision a certificat
J. Connect each device to AWS loT Core.
K. Set up an Amazon API Gateway HTTP API and a Network Load Balancer (NLB). Create integration between API Gateway and the NL
L. Configure a mutual TLS certificate authorizer on the HTTP AP
M. Run an MQTT broker on an Amazon EC2 instance that the NLB target
N. Connect each device to the NLB.

Answer: D

Explanation:
This solution requires minimal operational overhead, as it only requires setting up AWS IoT Core and creating a thing for each device. (Reference: AWS Certified Solutions Architect - Professional Official Amazon Text Book, Page 537)
AWS IoT Core is a fully managed service that enables secure, bi-directional communication between internet-connected devices and the AWS Cloud. It supports the MQTT protocol and includes built-in device
authentication and access control. By using AWS IoT Core, the company can easily provision and manage the X.509 certificates for each device, and connect the devices to the service with minimal operational overhead.

NEW QUESTION 10
- (Exam Topic 1)
A company is migrating some of its applications to AWS. The company wants to migrate and modernize the applications quickly after it finalizes networking and security strategies. The company has set up an AWS Direct Connection connection in a central network account.
The company expects to have hundreds of AWS accounts and VPCs in the near future. The corporate network must be able to access the resources on AWS seamlessly and also must be able to communicate with all the VPCs. The company also wants to route its cloud resources to the internet through its on-premises data center.
Which combination of steps will meet these requirements? (Choose three.)

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

Answer: BDF

Explanation:

> Option A is incorrect because creating a Direct Connect gateway in the central account and creating an association proposal by using the Direct Connect gateway and the account ID for every virtual private gateway does not enable active-passive failover between the regions. A Direct Connect gateway is a globally available resource that enables you to connect your AWS Direct Connect connection over a private virtual interface (VIF) to one or more VPCs in any AWS Region. A virtual private gateway is the VPN concentrator on the Amazon side of a VPN connection. You can associate a Direct Connect gateway with either a transit gateway or a virtual private gateway. However, a Direct Connect gateway does not provide any load balancing or failover capabilities by itself1

> Option B is correct because creating a Direct Connect gateway and a transit gateway in the central network account and attaching the transit gateway to the Direct Connect gateway by using a transit VIF meets the requirement of enabling the corporate network to access the resources on AWS seamlessly and also to communicate with all the VPCs. A transit VIF is a type of private VIF that you can use to connect your AWS Direct Connect connection to a transit gateway or a Direct Connect gateway. A transit gateway is a network transit hub that you can use to interconnect your VPCs and on-premises
networks. By using a transit VIF, you can route traffic between your on-premises network and multiple
VPCs across different AWS accounts and Regions through a single connection23

> Option C is incorrect because provisioning an internet gateway, attaching the internet gateway to subnets, and allowing internet traffic through the gateway does not meet the requirement of routing cloud resources to the internet through its on-premises data center. An internet gateway is a horizontally scaled, redundant, and highly available VPC component that allows communication between your VPC and the internet. An internet gateway serves two purposes: to provide a target in your VPC route tables for internet-routable traffic, and to perform network address translation (NAT) for instances that have been assigned public IPv4 addresses. By using an internet gateway, you are routing cloud resources directly to the internet, not through your on-premises data center.

> Option D is correct because sharing the transit gateway with other accounts and attaching VPCs to the transit gateway meets the requirement of enabling the corporate network to access the resources on AWS seamlessly and also to communicate with all the VPCs. You can share your transit gateway with other AWS accounts within the same organization by using AWS Resource Access Manager (AWS RAM). This allows you to centrally manage connectivity from multiple accounts without having to create individual peering connections between VPCs or duplicate network appliances in each account. You can attach VPCs from different accounts and Regions to your shared transit gateway and enable routing between them.

> Option E is incorrect because provisioning VPC peering as necessary does not meet the requirement of enabling the corporate network to access the resources on AWS seamlessly and also to communicate with all the VPCs. VPC peering is a networking connection between two VPCs that enables you to route traffic between them using private IPv4 addresses or IPv6 addresses. You can create a VPC peering connection between your own VPCs, or with a VPC in another AWS account within a single Region. However, VPC peering does not allow you to route traffic from your on-premises network to your VPCs or between multiple Regions. You would need to create multiple VPN connections or Direct Connect connections for each VPC peering connection, which increases operational complexity and costs.

> Option F is correct because provisioning only private subnets, opening the necessary route on the transit gateway and customer gateway to allow outbound internet traffic from AWS to flow through NAT services that run in the data center meets the requirement of routing cloud resources to the internet through its on-premises data center. A private subnet is a subnet that's associated with a route table that has no route to an internet gateway. Instances in a private subnet can communicate with other instances in the same VPC but cannot access resources on the internet directly. To enable outbound internet access from instances in private subnets, you can use NAT devices such as NAT gateways or NAT instances that are deployed in public subnets. A public subnet is a subnet that's associated with a route table that has a route to an internet gateway. Alternatively, you can use your on-premises data center as a NAT device by configuring routes on your transit gateway and customer gateway that direct outbound internet traffic from your private subnets through your VPN connection or Direct Connect connection. This way, you can route cloud resources to the internet through your on-premises data center instead of using an internet gateway.
References: 1:
https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-gateways-intro.html 2:
https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-transit-virtual-interfaces.html 3: https://docs.aws.amazon.com/vpc/latest/tgw/what-is-transit-gateway.html : https://docs.aws.amazon.com/vpc/latest/userguide/VPC_Internet_Gateway.html : https://docs.aws.amazon.com/vpc/latest/tgw/tgw-sharing.html : https://docs.aws.amazon.com/vpc/latest/peering/what-is-vpc-peering.html : https://docs.aws.amazon.com/vpc/latest/userguide/VPC_Scenario2.html : https://docs.aws.amazon.com/vpc/latest/userguide/VPC_Scenario3.html : https://docs.aws.amazon.com/vpc/latest/userguide/VPC_NAT_Instance.html : https://docs.aws.amazon.com/vpc/latest/userguide/VPC_NAT_Gateway.html

NEW QUESTION 10
- (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 12
- (Exam Topic 1)
A company is running an application in the AWS Cloud. The application runs on containers in an Amazon Elastic Container Service (Amazon ECS) cluster. The ECS tasks use the Fargate launch type. The application's data is relational and is stored in Amazon Aurora MySQL. To meet regulatory requirements, the application must be able to recover to a separate AWS Region in the event of an application failure. In case of a failure, no data can be lost. Which solution will meet these requirements with the LEAST amount of operational overhead?

A. Provision an Aurora Replica in a different Region.
B. Set up AWS DataSync for continuous replication of the data to a different Region.
C. Set up AWS Database Migration Service (AWS DMS) to perform a continuous replication of the data to a different Region.
D. Use Amazon Data Lifecycle Manager {Amazon DLM) to schedule a snapshot every 5 minutes.

Answer: A

Explanation:
Provision an Aurora Replica in a different Region will meet the requirement of the application being able to recover to a separate AWS Region in the event of an application failure, and no data can be lost, with the least amount of operational overhead.

NEW QUESTION 16
- (Exam Topic 1)
A software company has deployed an application that consumes a REST API by using Amazon API Gateway. AWS Lambda functions. and an Amazon DynamoDB table. The application is showing an increase in the number of errors during PUT requests. Most of the PUT calls come from a small number of clients that are authenticated with specific API keys.
A solutions architect has identified that a large number of the PUT requests originate from one client. The API is noncritical, and clients can tolerate retries of unsuccessful calls. However, the errors are displayed to customers and are causing damage to the API's reputation.
What should the solutions architect recommend to improve the customer experience?

A. Implement retry logic with exponential backoff and irregular variation in the client applicatio
B. Ensure that the errors are caught and handled with descriptive error messages.
C. Implement API throttling through a usage plan at the API Gateway leve
D. Ensure that the client application handles code 429 replies without error.
E. Turn on API caching to enhance responsiveness for the production stag
F. Run 10-minute load tests.Verify that the cache capacity is appropriate for the workload.
G. Implement reserved concurrency at the Lambda function level to provide the resources that are needed during sudden increases in traffic.

Answer: B

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/aws-batch-requests-error/ https://aws.amazon.com/premiumsupport/knowledge-center/api-gateway-429-limit/

NEW QUESTION 17
- (Exam Topic 1)
A company is running an application on several Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer. The load on the application varies throughout the day, and EC2 instances are scaled in and out on a regular basis. Log files from the EC2 instances are copied to a central Amazon S3 bucket every 15 minutes. The security team discovers that log files are missing from some of the terminated EC2 instances.
Which set of actions will ensure that log files are copied to the central S3 bucket from the terminated EC2 instances?

A. Create a script to copy log files to Amazon S3, and store the script in a file on the EC2 instanc
B. Create an Auto Scaling lifecycle hook and an Amazon EventBridge (Amazon CloudWatch Events) rule to detect lifecycle events from the Auto Scaling grou
C. Invoke an AWS Lambda function on the autoscaling:EC2_INSTANCE_TERMINATING transition to send ABANDON to the Auto Scaling group to prevent termination, run the script to copy the log files, and terminate the instance using the AWS SDK.
D. Create an AWS Systems Manager document with a script to copy log files to Amazon S3. Create an Auto Scaling lifecycle hook and an Amazon EventBridge (Amazon CloudWatch Events) rule to detect lifecycle events from the Auto Scaling grou
E. Invoke an AWS Lambda function on the autoscaling:EC2_INSTANCE_TERMINATING transition to call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send CONTINUE to the Auto Scaling group to terminate the instance.
F. Change the log delivery rate to every 5 minute

G. Create a script to copy log files to Amazon S3, and add the script to EC2 instance user dat
H. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to detect EC2 instance terminatio
I. Invoke an AWS Lambda function from the EventBridge (CloudWatch Events) rule that uses the AWS CLI to run the user-data script to copy the log files and terminate the instance.
J. Create an AWS Systems Manager document with a script to copy log files to Amazon S3. Create an Auto Scaling lifecycle hook that publishes a message to an Amazon Simple Notification Service (Amazon SNS) topi
K. From the SNS notification, call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send ABANDON to the Auto Scaling group to terminate the instance.

Answer: B

Explanation:
https://docs.aws.amazon.com/autoscaling/ec2/userguide/adding-lifecycle-hooks.html
- Refer to Default Result section - If the instance is terminating, both abandon and continue allow the instance to terminate. However, abandon stops any remaining actions, such as other lifecycle hooks, and continue allows any other lifecycle hooks to complete.
https://aws.amazon.com/blogs/infrastructure-and-automation/run-code-before-terminating-an-ec2-auto-scaling-i https://github.com/aws-samples/aws-lambda-lifecycle-hooks-function
https://github.com/aws-samples/aws-lambda-lifecycle-hooks-function/blob/master/cloudformation/template.yam

NEW QUESTION 22
- (Exam Topic 1)
A finance company hosts a data lake in Amazon S3. The company receives financial data records over SFTP each night from several third parties. The company runs its own SFTP server on an Amazon EC2 instance in a public subnet of a VPC. After the files ate uploaded, they are moved to the data lake by a cron job that runs on the same instance. The SFTP server is reachable on DNS sftp.examWe.com through the use of Amazon Route 53.
What should a solutions architect do to improve the reliability and scalability of the SFTP solution?

A. Move the EC2 instance into an Auto Scaling grou
B. Place the EC2 instance behind an Application Load Balancer (ALB). Update the DNS record sftp.example.com in Route 53 to point to the ALB.
C. Migrate the SFTP server to AWS Transfer for SFT
D. Update the DNS record sftp.example.com in Route 53 to point to the server endpoint hostname.
E. Migrate the SFTP server to a file gateway in AWS Storage Gatewa
F. Update the DNS record sflp.example.com in Route 53 to point to the file gateway endpoint.
G. Place the EC2 instance behind a Network Load Balancer (NLB). Update the DNS record sftp.example.com in Route 53 to point to the NLB.

Answer: B

Explanation:
https://aws.amazon.com/aws-transfer-family/faqs/ https://docs.aws.amazon.com/transfer/latest/userguide/what-is-aws-transfer-family.html
https://aws.amazon.com/about-aws/whats-new/2018/11/aws-transfer-for-sftp-fully-managed-sftp-for-s3/?nc1=h_

NEW QUESTION 23
- (Exam Topic 1)
A company wants to migrate to AWS. The company wants to use a multi-account structure with centrally managed access to all accounts and applications. The company also wants to keep the traffic on a private network. Multi-factor authentication (MFA) is required at login, and specific roles are assigned to user groups. The company must create separate accounts for development. staging, production, and shared network. The production account and the shared network account must have connectivity to all accounts. The development account and the staging account must have access only to each other.
Which combination of steps should a solutions architect take 10 meet these requirements? (Choose three.)

A. Deploy a landing zone environment by using AWS Control Towe
B. Enroll accounts and invite existing accounts into the resulting organization in AWS Organizations.
C. Enable AWS Security Hub in all accounts to manage cross-account acces
D. Collect findings through AWS CloudTrail to force MFA login.
E. Create transit gateways and transit gateway VPC attachments in each accoun
F. Configure appropriate route tables.
G. Set up and enable AWS IAM Identity Center (AWS Single Sign-On). Create appropriate permission sets with required MFA for existing accounts.
H. Enable AWS Control Tower in all Recounts to manage routing between account
I. Collect findings through AWS CloudTrail to force MFA login.
J. Create IAM users and group
K. Configure MFA for all user
L. Set up Amazon Cognito user pools andidentity pools to manage access to accounts and between accounts.

Answer: ACD

Explanation:
The correct answer would be options A, C and D, because they address the requirements outlined in the question. A. Deploying a landing zone environment using AWS Control Tower and enrolling accounts in an organization in AWS Organizations allows for a centralized management of access to all accounts and applications. C. Creating transit gateways and transit gateway VPC attachments in each account and configuring appropriate route tables allows for private network traffic, and ensures that the production account and shared network account have connectivity to all accounts, while the development and staging accounts have access only to each other. D. Setting up and enabling AWS IAM Identity Center (AWS Single Sign-On) and creating appropriate permission sets with required MFA for existing accounts allows for multi-factor authentication at login and specific roles to be assigned to user groups.

NEW QUESTION 24
- (Exam Topic 1)
A solutions architect has developed a web application that uses an Amazon API Gateway Regional endpoint and an AWS Lambda function. The consumers of the web application are all close to the AWS Region where the application will be deployed. The Lambda function only queries an Amazon Aurora MySQL database.
The solutions architect has configured the database to have three read replicas.
During testing, the application does not meet performance requirements. Under high load, the application opens a large number of database connections. The solutions architect must improve the application's performance.
Which actions should the solutions architect take to meet these requirements? (Choose two.)

A. Use the cluster endpoint of the Aurora database.
B. Use RDS Proxy to set up a connection pool to the reader endpoint of the Aurora database.
C. Use the Lambda Provisioned Concurrency feature.
D. Move the code for opening the database connection in the Lambda function outside of the event handler.
E. Change the API Gateway endpoint to an edge-optimized endpoint.

Answer: BD

Explanation:
Connect to RDS outside of Lambda handler method to improve performance https://awstut.com/en/2022/04/30/connect-to-rds-outside-of-lambda-handler-method-to-improve-performance-en
Using RDS Proxy, you can handle unpredictable surges in database traffic. Otherwise, these surges might cause issues due to oversubscribing connections or creating new connections at a fast rate. RDS Proxy establishes a database connection pool and reuses connections in this pool. This approach avoids the memory and CPU overhead of opening a new database connection each time. To protect the database against oversubscription, you can control the number of database connections that are created. https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/rds-proxy.html

NEW QUESTION 29
- (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 31
- (Exam Topic 1)
A company wants to migrate an application to Amazon EC2 from VMware Infrastructure that runs in an on-premises data center. A solutions architect must preserve the software and configuration settings during the migration.
What should the solutions architect do to meet these requirements?

A. Configure the AWS DataSync agent to start replicating the data store to Amazon FSx for Windows FileServer Use the SMB share to host the VMware data stor
B. Use VM Import/Export to move the VMs to Amazon EC2.
C. Use the VMware vSphere client to export the application as an image in Open Virealization Format (OVF) format Create an Amazon S3 bucket to store the image in the destination AWS Regio
D. Create and apply an IAM role for VM Import Use the AWS CLI to run the EC2 import command.
E. . Configure AWS Storage Gateway for files service to export a Common Internet File System (CIFSJ shar
F. Create a backup copy to the shared folde
G. Sign in to the AWS Management Console and create an AMI from the backup copy Launch an EC2 instance that is based on the AMI.
H. Create a managed-instance activation for a hybrid environment in AWS Systems Manage
I. Download and install Systems Manager Agent on the on-premises VM Register the VM with Systems Manager to be a managed instance Use AWS Backup to create a snapshot of the VM and create an AM
J. Launch an EC2 instance that is based on the AMI

Answer: D

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

NEW QUESTION 33
- (Exam Topic 1)

A company gives users the ability to upload images from a custom application. The upload process invokes an AWS Lambda function that processes and stores the image in an Amazon S3 bucket. The application invokes the Lambda function by using a specific function version ARN.

The Lambda function accepts image processing parameters by using environment variables. The company often adjusts the environment variables of the Lambda function to achieve optimal image processing output. The company tests different parameters and publishes a new function version with the updated environment variables after validating results. This update process also requires frequent changes to the custom application to invoke the new function version ARN. These changes cause interruptions for users.

A solutions architect needs to simplify this process to minimize disruption to users. Which solution will meet these requirements with the LEAST operational overhead?

A. Directly modify the environment variables of the published Lambda function versio
B. Use theSLATEST version to test image processing parameters.
C. Create an Amazon DynamoDB table to store the image processing parameter
D. Modify the Lambda function to retrieve the image processing parameters from the DynamoDB table.
E. Directly code the image processing parameters within the Lambda function and remove the environment variable
F. Publish a new function version when the company updates the parameters.
G. Create a Lambda function alia
H. Modify the client application to use the function alias AR
I. Reconfigure the Lambda alias to point to new versions of the function when the company finishes testing.

Answer: D

Explanation:
A Lambda function alias allows you to point to a specific version of a function and also can be updated to point to a new version of the function without modifying the client application. This way, the company can test different versions of the function with different environment variables and, once the optimal parameters are found, update the alias to point to the new version, without the need to update the client application.

By using this approach, the company can simplify the process of updating the environment variables, minimize disruption to users, and reduce the operational overhead.
Reference:
AWS Lambda documentation: https://aws.amazon.com/lambda/
AWS Lambda Aliases documentation: https://docs.aws.amazon.com/lambda/latest/dg/aliases-intro.html AWS Lambda versioning and aliases documentation: https://aws.amazon.com/blogs/compute/versioning-aliases-in-aws-lambda/

NEW QUESTION 34
- (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 37
- (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 42
- (Exam Topic 1)
A company has developed a web application. The company is hosting the application on a group of Amazon EC2 instances behind an Application Load Balancer. The company wants to improve the security posture of the application and plans to use AWS WAF web ACLs. The solution must not adversely affect legitimate traffic to the application.
How should a solutions architect configure the web ACLs to meet these requirements?

A. Set the action of the web ACL rules to Coun
B. Enable AWS WAF logging Analyze the requests for false positives Modify the rules to avoid any false positive Over time change the action of the web ACL rules from Count to Block.
C. Use only rate-based rules in the web ACL
D. and set the throttle limit as high as possible Temporarily block all requests that exceed the limi
E. Define nested rules to narrow the scope of the rate tracking.
F. Set the action o' the web ACL rules to Bloc
G. Use only AWS managed rule groups in the web ACLs Evaluate the rule groups by using Amazon CloudWatch metrics with AWS WAF sampled requests or AWS WAF logs.
H. Use only custom rule groups in the web ACL
I. and set the action to Allow Enable AWS WAF logging Analyze the requests tor false positives Modify the rules to avoid any false positive Over time, change the action of the web ACL rules from Allow to Block.

Answer: A

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/waf-analyze-count-action-rules/

NEW QUESTION 44
- (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 45
- (Exam Topic 1)
A company that uses AWS Organizations allows developers to experiment on AWS. As part of the landing zone that the company has deployed, developers use their company email address to request an account. The company wants to ensure that developers are not launching costly services or running services unnecessarily. The company must give developers a fixed monthly budget to limit their AWS costs.
Which combination of steps will meet these requirements? (Choose three.)

A. Create an SCP to set a fixed monthly account usage limi
B. Apply the SCP to the developer accounts.
C. Use AWS Budgets to create a fixed monthly budget for each developer's account as part of the account creation process.
D. Create an SCP to deny access to costly services and component
E. Apply the SCP to the developer accounts.
F. Create an IAM policy to deny access to costly services and component
G. Apply the IAM policy to the developer accounts.
H. Create an AWS Budgets alert action to terminate services when the budgeted amount is reached.Configure the action to terminate all services.
I. Create an AWS Budgets alert action to send an Amazon Simple Notification Service (Amazon SNS) notification when the budgeted amount is reache
J. Invoke an AWS Lambda function to terminate all services.

Answer: BCF

Explanation:
> Option A is incorrect because creating an SCP to set a fixed monthly account usage limit is not possible.
SCPs are policies that specify the services and actions that users and roles can use in the member accounts of an AWS Organization. SCPs cannot enforce budget limits or prevent users from launching

costly services or running services unnecessarily1

≫ Option B is correct because using AWS Budgets to create a fixed monthly budget for each developer's account as part of the account creation process meets the requirement of giving developers a fixed monthly budget to limit their AWS costs. AWS Budgets allows you to plan your service usage, service costs, and instance reservations. You can create budgets that alert you when your costs or usage exceed (or are forecasted to exceed) your budgeted amount2

≫ Option C is correct because creating an SCP to deny access to costly services and components meets the requirement of ensuring that developers are not launching costly services or running services
unnecessarily. SCPs can restrict access to certain AWS services or actions based on conditions such as region, resource tags, or request time. For example, an SCP can deny access to Amazon Redshift clusters or Amazon EC2 instances with certain instance types1

≫ Option D is incorrect because creating an IAM policy to deny access to costly services and components is not sufficient to meet the requirement of ensuring that developers are not launching costly services or running services unnecessarily. IAM policies can only control access to resources within a single AWS account. If developers have multiple accounts or can create new accounts, they can bypass the IAM policy restrictions. SCPs can apply across multiple accounts within an AWS Organization and prevent users from creating new accounts that do not comply with the SCP rules3

≫ Option E is incorrect because creating an AWS Budgets alert action to terminate services when the budgeted amount is reached is not possible. AWS Budgets alert actions can only perform one of the following actions: apply an IAM policy, apply an SCP, or send a notification through Amazon SNS. AWS Budgets alert actions cannot terminate services directly.

≫ Option F is correct because creating an AWS Budgets alert action to send an Amazon SNS notification when the budgeted amount is reached and invoking an AWS Lambda function to terminate all services meets the requirement of giving developers a fixed monthly budget to limit their AWS costs. AWS Budgets alert actions can send notifications through Amazon SNS when a budget threshold is breached. Amazon SNS can trigger an AWS Lambda function that can perform custom logic such as terminating all services in the developer's account. This way, developers cannot exceed their budget limit and incur additional costs.
References: 1: https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html 2
: https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/budgets-create.html 3: https://docs.aws.amazon.com/IAM/latest/UserGuide/introduction.html :
https://docs.aws.amazon.com/cost-management/latest/userguide/budgets-actions.html : https://docs.aws.amazon.com/sns/latest/dg/sns-lambda.html :
https://docs.aws.amazon.com/lambda/latest/dg/welcome.html

NEW QUESTION 50
- (Exam Topic 1)
A company is in the process of implementing AWS Organizations to constrain its developers to use only Amazon EC2. Amazon S3 and Amazon DynamoDB. The developers account resides In a dedicated organizational unit (OU). The solutions architect has implemented the following SCP on the developers account:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowEC2",
            "Effect": "Allow",
            "Action": "ec2:*",
            "Resource": "*"
        },
        {
            "Sid": "AllowDynamoDB",
            "Effect": "Allow",
            "Action": "dynamodb:*",
            "Resource": "*"
        },
        {
            "Sid": "AllowS3",
            "Effect": "Allow",
            "Action": "s3:*",
            "Resource": "*"
        }
    ]
}
```

When this policy is deployed, IAM users in the developers account are still able to use AWS services that are not listed in the policy. What should the solutions architect do to eliminate the developers' ability to use services outside the scope of this policy?

A. Create an explicit deny statement for each AWS service that should be constrained
B. Remove the Full AWS Access SCP from the developer account's OU
C. Modify the Full AWS Access SCP to explicitly deny all services
D. Add an explicit deny statement using a wildcard to the end of the SCP

Answer: B

Explanation:
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_inheritance_auth.html

NEW QUESTION 54
- (Exam Topic 1)
A company recently acquired several other companies. Each company has a separate AWS account with a different billing and reporting method. The acquiring company has consolidated all the accounts into one organization in AWS Organizations. However, the acquiring company has found it difficult to generate a cost report that contains meaningful groups for all the teams.
The acquiring company's finance team needs a solution to report on costs for all the companies through a self-managed application.
Which solution will meet these requirements?

A. Create an AWS Cost and Usage Report for the organizatio

B. Define tags and cost categories in the repor
C. Create a table in Amazon Athen
D. Create an Amazon QuickSight dataset based on the Athena tabl
E. Share the dataset with the finance team.
F. Create an AWS Cost and Usage Report for the organizatio
G. Define tags and cost categories in the repor
H. Create a specialized template in AWS Cost Explorer that the finance department will use to build reports.
I. Create an Amazon QuickSight dataset that receives spending information from the AWS Price List Query AP
J. Share the dataset with the finance team.
K. Use the AWS Price List Query API to collect account spending informatio
L. Create a specialized template in AWS Cost Explorer that the finance department will use to build reports.

Answer: A

Explanation:
Creating an AWS Cost and Usage Report for the organization and defining tags and cost categories in the report will allow for detailed cost reporting for the different companies that have been consolidated into one organization. By creating a table in Amazon Athena and an Amazon QuickSight dataset based on the Athena table, the finance team will be able to easily query and generate reports on the costs for all the companies. The dataset can then be shared with the finance team for them to use for their reporting needs.

NEW QUESTION 55
- (Exam Topic 1)
A company is using AWS Organizations lo manage multiple AWS accounts For security purposes, the company requires the creation of an Amazon Simple Notification Service (Amazon SNS) topic that enables integration with a third-party alerting system in all the Organizations member accounts
A solutions architect used an AWS CloudFormation template to create the SNS topic and stack sets to automate the deployment of CloudFormation stacks Trusted access has been enabled in Organizations
What should the solutions architect do to deploy the CloudFormation StackSets in all AWS accounts?

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

Answer: C

Explanation:
https://aws.amazon.com/blogs/aws/use-cloudformation-stacksets-to-provision-resources-across-multiple-aws-ac

NEW QUESTION 56
- (Exam Topic 1)
A company is hosting a critical application on a single Amazon EC2 instance. The application uses an Amazon ElastiCache for Redis single-node cluster for an in-memory data store. The application uses an Amazon RDS for MariaDB DB instance for a relational database. For the application to function, each piece of the infrastructure must be healthy and must be in an active state.
A solutions architect needs to improve the application's architecture so that the infrastructure can automatically recover from failure with the least possible downtime.
Which combination of steps will meet these requirements? (Select THREE.)

A. Use an Elastic Load Balancer to distribute traffic across multiple EC2 instance
B. Ensure that the EC2 instances are part of an Auto Scaling group that has a minimum capacity of two instances.
C. Use an Elastic Load Balancer to distribute traffic across multiple EC2 instances Ensure that the EC2 instances are configured in unlimited mode.
D. Modify the DB instance to create a read replica in the same Availability Zon
E. Promote the read replica to be the primary DB instance in failure scenarios.
F. Modify the DB instance to create a Multi-AZ deployment that extends across two Availability Zones.
G. Create a replication group for the ElastiCache for Redis cluste
H. Configure the cluster to use an Auto Scaling group that has a minimum capacity of two instances.
I. Create a replication group for the ElastiCache for Redis cluste
J. Enable Multi-AZ on the cluster.

Answer: ADF

Explanation:
> Option A is correct because using an Elastic Load Balancer and an Auto Scaling group with a minimum capacity of two instances can improve the availability and scalability of the EC2 instances that host the application. The load balancer can distribute traffic across multiple instances and the Auto Scaling group can replace any unhealthy instances automatically1
> Option D is correct because modifying the DB instance to create a Multi-AZ deployment that extends across two Availability Zones can improve the availability and durability of the RDS for MariaDB
database. Multi-AZ deployments provide enhanced data protection and minimize downtime by automatically failing over to a standby replica in another Availability Zone in case of a planned or unplanned outage4
> Option F is correct because creating a replication group for the ElastiCache for Redis cluster and enabling Multi-AZ on the cluster can improve the availability

and fault tolerance of the in-memory data store. A replication group consists of a primary node and up to five read-only replica nodes that are synchronized with the primary node using asynchronous replication. Multi-AZ allows automatic failove to one of the replicas if the primary node fails or becomes unreachable6 References: 1:
https://docs.aws.amazon.com/elasticloadbalancing/latest/userguide/how-elastic-load-balancing-works.html 2:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/burstable-performance-instances-unlimited-mode.htm 3:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html 4:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZ.html 5:
https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/AutoScaling.html 6: https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/Replication.Redis.Groups.html

NEW QUESTION 58
- (Exam Topic 1)
A company is running a data-intensive application on AWS. The application runs on a cluster of hundreds of Amazon EC2 instances. A shared file system also runs on several EC2 instances that store 200 TB of data. The application reads and modifies the data on the shared file system and generates a report. The job runs once monthly, reads a subset of the files from the shared file system, and takes about 72 hours to complete. The compute instances scale in an Auto Scaling group, but the instances that host the shared file system run continuously. The compute and storage instances are all in the same AWS Region.
A solutions architect needs to reduce costs by replacing the shared file system instances. The file system must provide high performance access to the needed data for the duration of the 72-hour run.
Which solution will provide the LARGEST overall cost reduction while meeting these requirements?

A. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Intelligent-Tiering storage clas
B. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using lazy loadin
C. Use the new file system as the shared storage for the duration of the jo
D. Delete the file system when the job is complete.
E. Migrate the data from the existing shared file system to a large Amazon Elastic Block Store (Amazon EBS) volume with Multi-Attach enable
F. Attach the EBS volume to each of the instances by using a user data script in the Auto Scaling group launch template
G. Use the EBS volume as the shared storage for the duration of the jo
H. Detach the EBS volume when the job is complete.
I. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Standard storage clas
J. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using batch loadin
K. Use the new file system as the shared storage for the duration of the jo
L. Delete the file system when the job is complete.
M. Migrate the data from the existing shared file system to an Amazon S3 bucke
N. Before the job runs each month, use AWS Storage Gateway to create a file gateway with the data from Amazon S3. Use the file gateway as the shared storage for the jo
O. Delete the file gateway when the job is complete.

Answer: A

Explanation:
https://aws.amazon.com/blogs/storage/new-enhancements-for-moving-data-between-amazon-fsx-for-lustre-and

NEW QUESTION 60
- (Exam Topic 1)
A company that has multiple AWS accounts is using AWS Organizations. The company's AWS accounts host VPCs, Amazon EC2 instances, and containers. The company's compliance team has deployed a security tool in each VPC where the company has deployments. The security tools run on EC2 instances and send information to the AWS account that is dedicated for the compliance team. The company has tagged all the compliance-related resources with a key of "costCenter" and a value or "compliance".
The company wants to identify the cost of the security tools that are running on the EC2 instances so that the company can charge the compliance team's AWS account. The cost calculation must be as accurate as possible.
What should a solutions architect do to meet these requirements?

A. In the management account of the organization, activate the costCenter user-defined ta
B. Configure monthly AWS Cost and Usage Reports to save to an Amazon S3 bucket in the management accoun
C. Use the tag breakdown in the report to obtain the total cost for the costCenter tagged resources.
D. In the member accounts of the organization, activate the costCenter user-defined ta
E. Configure monthly AWS Cost and Usage Reports to save to an Amazon S3 bucket in the management accoun
F. Schedule a monthly AWS Lambda function to retrieve the reports and calculate the total cost for the costCenter tagged resources.
G. In the member accounts of the organization activate the costCenter user-defined ta
H. From the management account, schedule a monthly AWS Cost and Usage Repor
I. Use the tag breakdown in the report to calculate the total cost for the costCenter tagged resources.
J. Create a custom report in the organization view in AWS Trusted Adviso
K. Configure the report to generate a monthly billing summary for the costCenter tagged resources in the compliance team's AWS account.

Answer: A

Explanation:
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/custom-tags.html
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/configurecostallocreport.html

NEW QUESTION 62
- (Exam Topic 1)
A company has its cloud infrastructure on AWS A solutions architect needs to define the infrastructure as code. The infrastructure is currently deployed in one AWS Region. The company's business expansion plan includes deployments in multiple Regions across multiple AWS accounts
What should the solutions architect do to meet these requirements?

A. Use AWS CloudFormation templates Add IAM policies to control the various accounts Deploy the templates across the multiple Regions
B. Use AWS Organizations Deploy AWS CloudFormation templates from the management account Use AWS Control Tower to manage deployments across accounts
C. Use AWS Organizations and AWS CloudFormation StackSets Deploy a CloudFormation template from an account that has the necessary IAM permissions

D. Use nested stacks with AWS CloudFormation templates Change the Region by using nested stacks

Answer: C

Explanation:
https://aws.amazon.com/blogs/aws/new-use-aws-cloudformation-stacksets-for-multiple-accounts-in-an-aws-orga AWS Organizations allows the management of multiple AWS accounts as a single entity and AWS
CloudFormation StackSets allows creating, updating, and deleting stacks across multiple accounts and regions in an organization. This solution allows creating a single CloudFormation template that can be deployed across multiple accounts and regions, and also allows for the management of access and permissions for the different accounts through the use of IAM roles and policies in the management account.

NEW QUESTION 67
- (Exam Topic 1)
A company has a multi-tier web application that runs on a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are in an Auto Scaling group. The ALB and the Auto Scaling group are replicated in a backup AWS Region. The minimum value and the maximum value for the Auto Scaling group are set to zero. An Amazon RDS Multi-AZ DB instance stores the application's data. The DB instance has a read replica in the backup Region. The application presents an endpoint to end users by using an Amazon Route 53 record.
The company needs to reduce its RTO to less than 15 minutes by giving the application the ability to automatically fail over to the backup Region. The company does not have a large enough budget for an active-active strategy.
What should a solutions architect recommend to meet these requirements?

A. Reconfigure the application's Route 53 record with a latency-based routing policy that load balances traffic between the two ALB
B. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group value
C. Create an Amazon CloudWatch alarm that is based on the HTTPCode_Target_5XX_Count metric for the ALB in the primary Regio
D. Configure the CloudWatch alarm to invoke the Lambda function.
E. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group value
F. Configure Route 53 with a health check that monitors the web application and sends an Amazon Simple Notification Service (Amazon SNS) notification to the Lambda function when the health check status is unhealth
G. Update the application's Route 53 record with a failover policy that routes traffic to the ALB in the backup Region when a health check failure occurs.
H. Configure the Auto Scaling group in the backup Region to have the same values as the Auto Scaling group in the primary Regio
I. Reconfigure the application's Route 53 record with a latency-based routing policy that load balances traffic between the two ALB
J. Remove the read replic
K. Replace the read replica with a standalone RDS DB instanc
L. Configure Cross-Region Replication between the RDS DB instances by using snapshots and Amazon S3.
M. Configure an endpoint in AWS Global Accelerator with the two ALBs as equal weighted target
N. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group value
O. Create an Amazon CloudWatch alarm that is based on the HTTPCode_Target_5XX_Count metric for the ALB in the primary Regio
P. Configure the CloudWatch alarm to invoke the Lambda function.

Answer: B

Explanation:
an AWS Lambda function in the backup region to promote the read replica and modify the Auto Scaling group values, and then configuring Route 53 with a health check that monitors the web application and sends an Amazon SNS notification to the Lambda function when the health check status is unhealthy. Finally, the application's Route 53 record should be updated with a failover policy that routes traffic to the ALB in the backup region when a health check failure occurs. This approach provides automatic failover to the backup region when a health check failure occurs, reducing the RTO to less than 15 minutes. Additionally, this approach is cost-effective as it does not require an active-active strategy.

NEW QUESTION 68
- (Exam Topic 1)
A retail company has an on-premises data center in Europe. The company also has a multi-Region AWS presence that includes the eu-west-1 and us-east-1 Regions. The company wants to be able to route network traffic from its on-premises infrastructure into VPCs in either of those Regions. The company also needs to support traffic that is routed directly between VPCs in those Regions. No single points of failure can exist on the network.
The company already has created two 1 Gbps AWS Direct Connect connections from its on-premises data center. Each connection goes into a separate Direct Connect location in Europe for high availability. These two locations are named DX-A and DX-B, respectively. Each Region has a single AWS Transit Gateway that is configured to route all inter-VPC traffic within that Region.
Which solution will meet these requirements?

A. Create a private VIF from the DX-A connection into a Direct Connect gatewa
B. Create a private VIF from the DX-B connection into the same Direct Connect gateway for high availabilit
C. Associate both the eu-west-1 and us-east-1 transit gateways with the Direct Connect gatewa
D. Peer the transit gatewayswith each other to support cross-Region routing.
E. Create a transit VIF from the DX-A connection into a Direct Connect gatewa
F. Associate the eu-west-1 transit gateway with this Direct Connect gatewa
G. Create a transit VIF from the DX-B connection into a separate Direct Connect gatewa
H. Associate the us-east-1 transit gateway with this separate Direct Connect gateway
I. Peer the Direct Connect gateways with each other to support high availability and cross-Region routing.
J. Create a transit VIF from the DX-A connection into a Direct Connect gatewa
K. Create a transit VIF from the DX-B connection into the same Direct Connect gateway for high availabilit
L. Associate both the eu-west-1 and us-east-1 transit gateways with this Direct Connect gatewa
M. Configure the Direct Connect gateway to route traffic between the transit gateways.
N. Create a transit VIF from the DX-A connection into a Direct Connect gatewa
O. Create a transit VIF from the DX-B connection into the same Direct Connect gateway for high availabilit
P. Associate both the eu-west-1 and us-east-1 transit gateways with this Direct Connect gatewa
Q. Peer the transit gateways with each other to support cross-Region routing.

Answer: D

Explanation:
in this solution, two transit VIFs are created - one from the DX-A connection and one from the DX-B connection - into the same Direct Connect gateway for high availability. Both the eu-west-1 and us-east-1 transit gateways are then associated with this Direct Connect gateway. The transit gateways are then peered with

each other to support cross-Region routing. This solution meets the requirements of the company by creating a highly available connection between the on-premises data center and the VPCs in both the eu-west-1 and us-east-1 regions, and by enabling direct traffic routing between VPCs in those regions.

NEW QUESTION 73
- (Exam Topic 1)
A large company is running a popular web application. The application runs on several Amazon EC2 Linux Instances in an Auto Scaling group in a private subnet. An Application Load Balancer is targeting the Instances In the Auto Scaling group in the private subnet. AWS Systems Manager Session Manager Is configured, and AWS Systems Manager Agent is running on all the EC2 instances.
The company recently released a new version of the application Some EC2 instances are now being marked as unhealthy and are being terminated As a result, the application is running at reduced capacity A solutions architect tries to determine the root cause by analyzing Amazon CloudWatch logs that are collected from the application, but the logs are inconclusive
How should the solutions architect gain access to an EC2 instance to troubleshoot the issue1?

A. Suspend the Auto Scaling group's HealthCheck scaling proces
B. Use Session Manager to log in to an instance that is marked as unhealthy
C. Enable EC2 instance termination protection Use Session Manager to log In to an instance that is marked as unhealthy.
D. Set the termination policy to Oldestinstance on the Auto Scaling grou
E. Use Session Manager to log in to an instance that is marked as unhealthy
F. Suspend the Auto Scaling group's Terminate proces
G. Use Session Manager to log in to an instance thatis marked as unhealthy

Answer: D

Explanation:
https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-suspend-resume-processes.html

NEW QUESTION 78
- (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 82
- (Exam Topic 1)
A solutions architect is designing the data storage and retrieval architecture for a new application that a company will be launching soon. The application is designed to ingest millions of small records per minute from devices all around the world. Each record is less than 4 KB in size and needs to be stored in a durable location where it can be retrieved with low latency. The data is ephemeral and the company is required to store the data for 120 days only, after which the data can be deleted.
The solutions architect calculates that, during the course of a year, the storage requirements would be about 10-15 TB.
Which storage strategy is the MOST cost-effective and meets the design requirements?

A. Design the application to store each incoming record as a single .csv file in an Amazon S3 bucket to allow for indexed retrieva
B. Configure a lifecycle policy to delete data older than 120 days.
C. Design the application to store each incoming record in an Amazon DynamoDB table properly configured for the scal
D. Configure the DynamoOB Time to Live (TTL) feature to delete records older than 120 days.
E. Design the application to store each incoming record in a single table in an Amazon RDS MySQL databas
F. Run a nightly cron job that executes a query to delete any records older than 120 days.
G. Design the application to batch incoming records before writing them to an Amazon S3 bucke
H. Updatethe metadata for the object to contain the list of records in the batch and use the Amazon S3 metadata search feature to retrieve the dat
I. Configure a lifecycle policy to delete the data after 120 days.

Answer: B

Explanation:
DynamoDB with TTL, cheaper for sustained throughput of small items + suited for fast retrievals. S3 cheaper for storage only, much higher costs with writes. RDS not designed for this use case.

NEW QUESTION 85
- (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 86
- (Exam Topic 1)
A weather service provides high-resolution weather maps from a web application hosted on AWS in the
eu-west-1 Region. The weather maps are updated frequently and stored in Amazon S3 along with static HTML content. The web application is fronted by Amazon CloudFront.
The company recently expanded to serve users in the us-east-1 Region, and these new users report that viewing their respective weather maps is slow from time to time.
Which combination of steps will resolve the us-east-1 performance issues? (Choose two.)

A. Configure the AWS Global Accelerator endpoint for the S3 bucket in eu-west-1. Configure endpoint groups for TCP ports 80 and 443 in us-east-1.

B. Create a new S3 bucket in us-east-1. Configure S3 cross-Region replication to synchronize from the S3 bucket in eu-west-1.

C. Use Lambda@Edge to modify requests from North America to use the S3 Transfer Acceleration endpoint in us-east-1.

D. Use Lambda@Edge to modify requests from North America to use the S3 bucket in us-east-1.

E. Configure the AWS Global Accelerator endpoint for us-east-1 as an origin on the CloudFront distributio

F. Use Lambda@Edge to modify requests from North America to use the new origin.

Answer: BD

Explanation:
https://aws.amazon.com/about-aws/whats-new/2016/04/transfer-files-into-amazon-s3-up-to-300-percent-faster/

NEW QUESTION 90
- (Exam Topic 1)
A company has hundreds of AWS accounts. The company recently implemented a centralized internal process for purchasing new Reserved Instances and modifying existing Reserved Instances. This process requires all business units that want to purchase or modify Reserved Instances to submit requests to a dedicated team for procurement. Previously, business units directly purchased or modified Reserved Instances in their own respective AWS accounts autonomously.
A solutions architect needs to enforce the new process in the most secure way possible.
Which combination of steps should the solutions architect take to meet these requirements? (Choose two.)

A. Ensure that all AWS accounts are part of an organization in AWS Organizations with all features enabled.

B. Use AWS Config to report on the attachment of an IAM policy that denies access to the ec2:PurchaseReservedInstancesOffering action and the ec2:ModifyReservedInstances action.

C. In each AWS account, create an IAM policy that denies the ec2:PurchaseReservedInstancesOffering action and the ec2:ModifyReservedInstances action.

D. Create an SCP that denies the ec2:PurchaseReservedInstancesOffering action and the ec2:ModifyReservedInstances actio

E. Attach the SCP to each OU of the organization.

F. Ensure that all AWS accounts are part of an organization in AWS Organizations that uses the consolidated billing feature.

Answer: AD

Explanation:
All features – The default feature set that is available to AWS Organizations. It includes all the functionality of consolidated billing, plus advanced features that give you more control over accounts in your organization. For example, when all features are enabled the management account of the organization has full control over what member accounts can do. The management account can apply SCPs to restrict the services and actions that users (including the root user) and roles in an account can access. https://docs.aws.amazon.com/organizations/latest/userguide/orgs_getting-started_concepts.html#feature-set

NEW QUESTION 94
- (Exam Topic 1)
A company has a legacy monolithic application that is critical to the company's business. The company hosts the application on an Amazon EC2 instance that runs Amazon Linux 2. The company's application team receives a directive from the legal department to back up the data from the instance's encrypted Amazon Elastic Block Store (Amazon EBS) volume to an Amazon S3 bucket. The application team does not have the administrative SSH key pair for the instance. The application must continue to serve the users.
Which solution will meet these requirements?

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

Answer: C

Explanation:
Taking a snapshot of the EBS volume using Amazon Data Lifecycle Manager (DLM) will meet the requirements because it allows you to create a backup of the volume without the need to access the instance or its SSH key pair. Additionally, DLM allows you to schedule the backups to occur at specific intervals and also enables you to copy the snapshots to an S3 bucket. This approach will not impact the running application as the backup is performed on the EBS volume level.

NEW QUESTION 96
- (Exam Topic 1)
A company used Amazon EC2 instances to deploy a web fleet to host a blog site The EC2 instances are behind an Application Load Balancer (ALB) and are configured in an Auto ScaSng group The web application stores all blog content on an Amazon EFS volume.
The company recently added a feature 'or Moggers to add video to their posts, attracting 10 times the previous user traffic At peak times of day. users report buffering and timeout issues while attempting to reach the site or watch videos
Which is the MOST cost-efficient and scalable deployment that win resolve the issues for users?

A. Reconfigure Amazon EFS to enable maximum I/O.
B. Update the Nog site to use instance store volumes tor storag
C. Copy the site contents to the volumes atlaunch and to Amazon S3 al shutdown.
D. Configure an Amazon CloudFront distributio
E. Point the distribution to an S3 bucket, and migrate the videos from EFS to Amazon S3.
F. Set up an Amazon CloudFront distribution for all site contents, and point the distribution at the ALB.

Answer: C

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-https-connection-fails/ Using an Amazon S3 bucket
Using a MediaStore container or a MediaPackage channel Using an Application Load Balancer
Using a Lambda function URL
Using Amazon EC2 (or another custom origin)
Using CloudFront origin groups https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/restrict-access-to-load-balancer.html

NEW QUESTION 97
- (Exam Topic 1)
A company has many AWS accounts and uses AWS Organizations to manage all of them. A solutions architect must implement a solution that the company can use to share a common network across multiple accounts.
The company's infrastructure team has a dedicated infrastructure account that has a VPC. The infrastructure team must use this account to manage the network. Individual accounts cannot have the ability to manage their own networks. However, individual accounts must be able to create AWS resources within subnets.
Which combination of actions should the solutions architect perform to meet these requirements? (Select TWO.)

A. Create a transit gateway in the infrastructure account.
B. Enable resource sharing from the AWS Organizations management account.
C. Create VPCs in each AWS account within the organization in AWS Organization
D. Configure the VPCs to share the same CIDR range and subnets as the VPC in the infrastructure accoun
E. Peer the VPCs in each individual account with the VPC in the infrastructure account,
F. Create a resource share in AWS Resource Access Manager in the infrastructure accoun
G. Select the specific AWS Organizations OU that will use the shared networ
H. Select each subnet to associate with the resource share.
I. Create a resource share in AWS Resource Access Manager in the infrastructure accoun
J. Select the specific AWS Organizations OU that will use the shared networ
K. Select each prefix list to associate with the resource share.

Answer: AE

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

NEW QUESTION 98
- (Exam Topic 1)
A company has its cloud infrastructure on AWS A solutions architect needs to define the infrastructure as code. The infrastructure is currently deployed in one AWS Region. The company's business expansion plan includes deployments in multiple Regions across multiple AWS accounts
What should the solutions architect do to meet these requirements?

A. Use AWS CloudFormation templates Add IAM policies to control the various accounts Deploy the templates across the multiple Regions
B. Use AWS Organizations Deploy AWS CloudFormation templates from the management account Use AWS Control Tower to manage deployments across accounts
C. Use AWS Organizations and AWS CloudFormation StackSets Deploy a CloudFormation template from an account that has the necessary IAM permissions
D. Use nested stacks with AWS CloudFormation templates Change the Region by using nested stacks

Answer: C

Explanation:

https://aws.amazon.com/blogs/aws/new-use-aws-cloudformation-stacksets-for-multiple-accounts-in-an-aws-orga AWS Organizations allows the management of multiple AWS accounts as a single entity and AWS

CloudFormation StackSets allows creating, updating, and deleting stacks across multiple accounts and regions in an organization. This solution allows creating a single CloudFormation template that can be deployed across multiple accounts and regions, and also allows for the management of access and permissions for the different accounts through the use of IAM roles and policies in the management account.

NEW QUESTION 99
- (Exam Topic 1)
A company wants to deploy an AWS WAF solution to manage AWS WAF rules across multiple AWS accounts. The accounts are managed under different OUs in AWS Organizations.
Administrators must be able to add or remove accounts or OUs from managed AWS WAF rule sets as needed Administrators also must have the ability to automatically update and remediate noncompliant AWS WAF rules in all accounts
Which solution meets these requirements with the LEAST amount of operational overhead?

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

Answer: A

Explanation:

https://aws.amazon.com/solutions/implementations/automations-for-aws-firewall-manager/
In this solution, AWS Firewall Manager is used to manage AWS WAF rules across accounts in the organization. An AWS Systems Manager Parameter Store parameter is used to store account numbers and OUs to manage. This parameter can be updated as needed to add or remove accounts or OUs. An Amazon EventBridge rule is used to identify any changes to the parameter and to invoke an AWS Lambda function to update the security policy in the Firewall Manager administrative account. This solution allows for easy management of AWS WAF rules across multiple accounts with minimal operational overhead

NEW QUESTION 104
- (Exam Topic 1)
A company runs a proprietary stateless ETL application on an Amazon EC2 Linux instance. The application is a Linux binary, and the source code cannot be modified. The application is single-threaded, uses 2 GB of RAM. and is highly CPU intensive The application is scheduled to run every 4 hours and runs for up to 20 minutes A solutions architect wants to revise the architecture for the solution.
Which strategy should the solutions architect use?

A. Use AWS Lambda to run the applicatio
B. Use Amazon CloudWatch Logs to invoke the Lambda function every 4 hours.
C. Use AWS Batch to run the applicatio
D. Use an AWS Step Functions state machine to invoke the AWS Batch job every 4 hours.
E. Use AWS Fargate to run the applicatio
F. Use Amazon EventBridge (Amazon CloudWatch Events) to invoke the Fargate task every 4 hours.
G. Use Amazon EC2 Spot Instances to run the applicatio
H. Use AWS CodeDeploy to deploy and run the application every 4 hours.

Answer: C

Explanation:

step function could run a scheduled task when triggered by eventbrige, but why would you add that layer of complexity just to run aws batch when you could directly invoke it through eventbridge. The link provided - https://aws.amazon.com/pt/blogs/compute/orchestrating-high-performance-computing-with-aws-step-functions- makes sense only for HPC, this is a single instance that needs to be run

NEW QUESTION 105
- (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 107
- (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 108
- (Exam Topic 1)
A company has an on-premises monitoring solution using a PostgreSQL database for persistence of events. The database is unable to scale due to heavy ingestion and it frequently runs out of storage.
The company wants to create a hybrid solution and has already set up a VPN connection between its network and AWS. The solution should include the following attributes:
• Managed AWS services to minimize operational complexity
• A buffer that automatically scales to match the throughput of data and requires no on-going administration.
• A visualization toot to create dashboards to observe events in near-real time.
• Support for semi -structured JSON data and dynamic schemas.
Which combination of components will enabled© company to create a monitoring solution that will satisfy these requirements'' (Select TWO.)

A. Use Amazon Kinesis Data Firehose to buffer events Create an AWS Lambda function 10 process and transform events
B. Create an Amazon Kinesis data stream to buffer events Create an AWS Lambda function to process and transform evens
C. Configure an Amazon Aurora PostgreSQL DB cluster to receive events Use Amazon Quick Sight to read from the database and create near-real-time visualizations and dashboards
D. Configure Amazon Elasticsearch Service (Amazon ES) to receive events Use the Kibana endpoint deployed with Amazon ES to create near-real-time visualizations and dashboards.
E. Configure an Amazon Neptune 0 DB instance to receive events Use Amazon QuickSight to read from the database and create near-real-time visualizations and dashboards

Answer: AD

Explanation:
https://aws.amazon.com/kinesis/data-firehose/faqs/

NEW QUESTION 113
- (Exam Topic 2)
A company's solutions architect is analyzing costs of a multi-application environment. The environment is deployed across multiple Availability Zones in a single AWS Region. After a recent acquisition, the company manages two organizations in AWS Organizations. The company has created multiple service provider applications as AWS PrivateLink-powered VPC endpoint services in one organization. The company has created multiple service consumer applications in the other organization.
Data transfer charges are much higher than the company expected, and the solutions architect needs to reduce the costs. The solutions architect must recommend guidelines for developers to follow when they deploy services. These guidelines must minimize data transfer charges for the whole environment.
Which guidelines meet these requirements? (Select TWO.)

A. Use AWS Resource Access Manager to share the subnets that host the service provider applications with other accounts in the organization.
B. Place the service provider applications and the service consumer applications in AWS accounts in the same organization.
C. Turn off cross-zone load balancing for the Network Load Balancer in all service provider application deployments.
D. Ensure that service consumer compute resources use the Availability Zone-specific endpoint service by using the endpoint's local DNS name.

E. Create a Savings Plan that provides adequate coverage for the organization's planned inter-Availability Zone data transfer usage.

Answer: CD

Explanation:
Cross-zone load balancing enables traffic to be distributed evenly across all registered instances in all enabled Availability Zones. However, this also increases data transfer charges between Availability Zones. By turning off cross-zone load balancing, the service provider applications can reduce inter-Availability Zone data transfer costs. Similarly, by using the Availability Zone-specific endpoint service, the service consumer applications can ensure that they connect to the nearest service provider application in the same Availability Zone, avoiding cross-Availability Zone data transfer charges. References:

≫ https://docs.aws.amazon.com/vpc/latest/userguide/vpce-interface.html#vpce-interface-dns

NEW QUESTION 115
- (Exam Topic 2)
A company has a website that runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are in an Auto Scaling group. The ALB is associated with an AWS WAF web ACL.
The website often encounters attacks in the application layer. The attacks produce sudden and significant increases in traffic on the application server. The access logs show that each attack originates from different IP addresses. A solutions architect needs to implement a solution to mitigate these attacks.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create an Amazon CloudWatch alarm that monitors server acces
B. Set a threshold based on access by IP addres
C. Configure an alarm action that adds the IP address to the web ACL's deny list.
D. Deploy AWS Shield Advanced in addition to AWS WA
E. Add the ALB as a protected resource.
F. Create an Amazon CloudWatch alarm that monitors user IP addresse
G. Set a threshold based on access by IP addres
H. Configure the alarm to invoke an AWS Lambda function to add a deny rule in the application server's subnet route table for any IP addresses that activate the alarm.
I. Inspect access logs to find a pattern of IP addresses that launched the attack
J. Use an Amazon Route 53 geolocation routing policy to deny traffic from the countries that host those IP addresses.

Answer: C

Explanation:
"The AWS WAF API supports security automation such as blacklisting IP addresses that exceed request limits, which can be useful for mitigating HTTP flood attacks." >
https://aws.amazon.com/blogs/security/how-to-protect-dynamic-web-applications-against-ddos-attacks-by-using

NEW QUESTION 116
- (Exam Topic 2)
A company uses AWS Organizations for a multi-account setup in the AWS Cloud. The company's finance team has a data processing application that uses AWS Lambda and Amazon DynamoDB. The company's marketing team wants to access the data that is stored in the DynamoDB table.
The DynamoDB table contains confidential data. The marketing team can have access to only specific attributes of data in the DynamoDB table. The fi-nance team and the marketing team have separate AWS accounts.
What should a solutions architect do to provide the marketing team with the appropriate access to the DynamoDB table?

A. Create an SCP to grant the marketing team's AWS account access to the specific attributes of the DynamoDB tabl
B. Attach the SCP to the OU of the finance team.
C. Create an IAM role in the finance team's account by using IAM policy conditions for specific DynamoDB attributes (fine-grained access con-trol). Establish trust with the marketing team's accoun
D. In the mar-keting team's account, create an IAM role that has permissions to as-sume the IAM role in the finance team's account.
E. Create a resource-based IAM policy that includes conditions for spe-cific DynamoDB attributes (fine-grained access control). Attach the policy to the DynamoDB tabl
F. In the marketing team'saccount, create an IAM role that has permissions to access the DynamoDB table in the finance team's account.
G. Create an IAM role in the finance team's account to access the Dyna-moDB tabl
H. Use an IAM permissions boundary to limit the access to the specific attribute
I. In the marketing team's account, create an IAM role that has permissions to assume the IAM role in the finance team's account.

Answer: C

Explanation:
The company should create a resource-based IAM policy that includes conditions for specific DynamoDB attributes (fine-grained access control). The company should attach the policy to the DynamoDB table. In the marketing team's account, the company should create an IAM role that has permissions to access the DynamoDB table in the finance team's account. This solution will meet the requirements because a
resource-based IAM policy is a policy that you attach to an AWS resource (such as a DynamoDB table) to control who can access that resource and what actions they can perform on it. You can use IAM policy conditions to specify fine-grained access control for DynamoDB items and attributes. For example, you can allow or deny access to specific attributes of all items in a table by matching on attribute names1. By creating a resource-based policy that allows access to only specific attributes of the DynamoDB table and attaching it to the table, the company can restrict access to confidential data. By creating an IAM role in the marketing team's account that has permissions to access the DynamoDB table in the finance team's account, the company can enable cross-account access.
The other options are not correct because:

≫ Creating an SCP to grant the marketing team's AWS account access to the specific attributes of the DynamoDB table would not work because SCPs are policies that you can use with AWS Organizations to manage permissions in your organization's accounts. SCPs do not grant permissions; instead, they specify the maximum permissions that identities in an account can have2. SCPs cannot be used to specify fine-grained access control for DynamoDB items and attributes.

≫ Creating an IAM role in the finance team's account by using IAM policy conditions for specific DynamoDB attributes and establishing trust with the marketing team's account would not work because IAM roles are identities that you can create in your account that have specific permissions. You can use an IAM role to delegate access to users, applications, or services that don't normally have access to your AWS resources3. However, creating an IAM role in the finance team's account would not restrict access to specific attributes of the DynamoDB table; it would only allow cross-account access. The company would still need a resource-based policy attached to the table to enforce fine-grained access control.

≫ Creating an IAM role in the finance team's account to access the DynamoDB table and using an IAM permissions boundary to limit the access to the specific

attributes would not work because IAM permissions boundaries are policies that you use to delegate permissions management to other
users. You can use permissions boundaries to limit the maximum permissions that an identity-based
policy can grant to an IAM entity (user or role)4. Permissions boundaries cannot be used to specify fine-grained access control for DynamoDB items and attributes.
References:

> https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/specifying-conditions.html

> https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html

> https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html

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

NEW QUESTION 121
- (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 126
- (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 131
- (Exam Topic 2)
A company needs to establish a connection from its on-premises data center to AWS. The company needs to connect all of its VPCs that are located in different AWS Regions with transitive routing capabilities between VPC networks. The company also must reduce network outbound traffic costs, increase bandwidth throughput, and provide a consistent network experience for end users.
Which solution will meet these requirements?

A. Create an AWS Site-to-Site VPN connection between the on-premises data center and a new central VP
B. Create VPC peering connections that initiate from the central VPC to all other VPCs.
C. Create an AWS Direct Connect connection between the on-premises data center and AW
D. Provision a transit VIF, and connect it to a Direct Connect gatewa
E. Connect the Direct Connect gateway to all the other VPCs by using a transit gateway in each Region.
F. Create an AWS Site-to-Site VPN connection between the on-premises data center and a new central VP
G. Use a transit gateway with dynamic routin
H. Connect the transit gateway to all other VPCs.
I. Create an AWS Direct Connect connection between the on-premises data center and AWS Establish an AWS Site-to-Site VPN connection between all VPCs in each Regio
J. Create VPC peering connections that initiate from the central VPC to all other VPCs.

Answer: B

Explanation:
Transit GW + Direct Connect GW + Transit VIF + enabled SiteLink if two different DX locations https://aws.amazon.com/blogs/networking-and-content-delivery/introducing-aws-direct-connect-sitelink/

NEW QUESTION 134
......

Thank You for Trying Our Product

* 100% Pass or Money Back

    All our products come with a 90-day Money Back Guarantee.

* One year free update

    You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

    We currently serve more than 30,000,000 customers.

* Shop Securely

    All transactions are protected by VeriSign!

100% Pass Your SAP-C02 Exam with Our Prep Materials Via below:

https://www.certleader.com/SAP-C02-dumps.html