CEplus

Number: SAP-C02
Passing Score: 800
Time Limit: 120
File Version: 7.0

Exam Code: SAP-C02

Exam Name: AWS Certified Solutions Architect - Professional

VCEûp CEplus

CEplus VCEûp

VCEplus

Exam A

QUESTION 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 bucket. Configure the S3 bucket to use the S3 One Zone-Infrequent Access (S3 One Zone-IA) storage class as default. Configure the S3 bucket for website hosting. Create an S3 interface endpoint. Configure the S3 bucket to allow access only through that endpoint.
B. Launch an Amazon EC2 instance that runs a web server. 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.
C. Launch an Amazon EC2 instance that runs a web server Attach an Amazon Elastic Block Store (Amazon EBS) volume to store the archived data. Use the Cold HDD (sc1) volume type. Configure the instance security groups to allow access only from private networks.
D. Create an Amazon S3 bucket. Configure the S3 bucket to use the S3 Glacier Deep Archive storage class as default. Configure the S3 bucket for website hosting. Create an S3 interface endpoint. Configure the S3 bucket to allow access only through that endpoint.

Correct Answer: D
Section:
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.

QUESTION 2
A company is using an on-premises Active Directory service for user authentication. The company wants to use the same authentication service to sign in to the company's AWS accounts, which are using AWS Organizations. AWS Site-to-Site VPN connectivity already exists between the on-premises environment and all the company's AWS accounts.
The company's security policy requires conditional access to the accounts based on user groups and roles. User identities must be managed in a single location.
Which solution will meet these requirements?

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

Correct Answer: D
Section:
Explanation:

VCEûp VCEplus IT Certification Exams - Questions & Answers | VCEplus.io VCEplus VCEûp

CEplus

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

QUESTION 3
A software company has deployed an application that consumes a REST API by using Amazon API Gateway. AWS Lambda functions, and an Amazon DynamoDB table. The application is showing an increase in the number of errors during PUT requests. Most of the PUT calls come from a small number of clients that are authenticated with specific API keys.
A solutions architect has identified that a large number of the PUT requests originate from one client. The API is noncritical, and clients can tolerate retries of unsuccessful calls. However, the errors are displayed to customers and are causing damage to the API's reputation.
What should the solutions architect recommend to improve the customer experience?

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

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

QUESTION 4
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 dat

A. 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?
B. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Intelligent-Tiering storage class. 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 loading. Use the new file system as the shared storage for the duration of the job. Delete the file system when the job is complete.
C. Migrate the data from the existing shared file system to a large Amazon Elastic Block Store (Amazon EBS) volume with Multi-Attach enabled. Attach the EBS volume to each of the instances by using a user data script in the Auto Scaling group launch template. Use the EBS volume as the shared storage for the duration of the job. Detach the EBS volume when the job is complete.
D. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Standard storage class. 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 loading. Use the new file system as the shared storage for the duration of the job. Delete the file system when the job is complete.
E. Migrate the data from the existing shared file system to an Amazon S3 bucket. 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 job. Delete the file gateway when the job is complete.

Correct Answer: A
Section:
Explanation:
https://aws.amazon.com/blogs/storage/new-enhancements-for-moving-data-between-amazon-fsx-for-lustre-and-amazon-s3/

QUESTION 5
A company uses a service to collect metadata from applications that the company hosts on premises. Consumer devices such as TVs and internet radios access the applications. Many older devices do not support certain HTTP headers and exhibit errors when these headers are present in responses. The company has configured an on-premises load balancer to remove the unsupported headers from responses sent to older devices, which the company identified by the User-Agent headers.
The company wants to migrate the service to AWS, adopt serverless technologies, and retain the ability to support the older devices. The company has already migrated the applications into a set of AWS Lambda functions.

VCEûp CEplus IT Certification Exams - Questions & Answers | VCEplus.io CEplus VCEûp

VCEplus

Which solution will meet these requirements?

A. Create an Amazon CloudFront distribution for the metadata service. Create an Application Load Balancer (ALB). Configure the CloudFront distribution to forward requests to the ALB. Configure the ALB to invoke the correct Lambda function for each type of request. Create a CloudFront function to remove the problematic headers based on the value of the User-Agent header.
B. Create an Amazon API Gateway REST API for the metadata service. Configure API Gateway to invoke the correct Lambda function for each type of request. Modify the default gateway responses to remove the problematic headers based on the value of the User-Agent header.
C. Create an Amazon API Gateway HTTP API for the metadata service. Configure API Gateway to invoke the correct Lambda function for each type of request. Create a response mapping template to remove the problematic headers based on the value of the User-Agent. Associate the response data mapping with the HTTP API.
D. Create an Amazon CloudFront distribution for the metadata service. Create an Application Load Balancer (ALB). Configure the CloudFront distribution to forward requests to the ALB. Configure the ALB to invoke the correct Lambda function for each type of request. Create a Lambda@Edge function that will remove the problematic headers in response to viewer requests based on the value of the User-Agent header.

Correct Answer: D
Section:
Explanation:
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html

QUESTION 6
A company is running a traditional web application on Amazon EC2 instances. The company needs to refactor the application as microservices that run on containers. Separate versions of the application exist in two distinct environments: production and testing. Load for the application is variable, but the minimum load and the maximum load are known. A solutions architect needs to design the updated application with a serverless architecture that minimizes operational complexity.
Which solution will meet these requirements MOST cost-effectively?

A. Upload the container images to AWS Lambda as functions. Configure a concurrency limit for the associated Lambda functions to handle the expected peak load. Configure two separate Lambda integrations within Amazon API Gateway: one for production and one for testing.
B. Upload the container images to Amazon Elastic Container Registry (Amazon ECR). Configure two auto scaled Amazon Elastic Container Service (Amazon ECS) clusters with the Fargate launch type to handle the expected load. Deploy tasks from the ECR images. Configure two separate Application Load Balancers to direct traffic to the ECS clusters.
C. Upload the container images to Amazon Elastic Container Registry (Amazon ECR). Configure two auto scaled Amazon Elastic Kubernetes Service (Amazon EKS) clusters with the Fargate launch type to handle the expected load. Deploy tasks from the ECR images. Configure two separate Application Load Balancers to direct traffic to the EKS clusters.
D. Upload the container images to AWS Elastic Beanstalk. In Elastic Beanstalk, create separate environments and deployments for production and testing. Configure two separate Application Load Balancers to direct traffic to the Elastic Beanstalk deployments.

Correct Answer: D
Section:
Explanation:
minimizes operational + microservices that run on containers = AWS Elastic Beanstalk

QUESTION 7
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 dat

A. 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?
B. Reconfigure the application's Route 53 record with a latency-based routing policy that load balances traffic between the two ALBs. Create an AWS Lambda function in the backup Region to promote the read

VCEûp VCEplus IT Certification Exams - Questions & Answers | VCEplus.io VCEplus VCEûp

VCEplus

replica and modify the Auto Scaling group values. Create an Amazon CloudWatch alarm that is based on the HTTPCode_Target_5XX_Count metric for the ALB in the primary Region. Configure the CloudWatch alarm to invoke the Lambda function.

C. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group values. 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 unhealthy. 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.

D. Configure the Auto Scaling group in the backup Region to have the same values as the Auto Scaling group in the primary Region. Reconfigure the application's Route 53 record with a latency-based routing policy that load balances traffic between the two ALBs. Remove the read replica. Replace the read replica with a standalone RDS DB instance. Configure Cross-Region Replication between the RDS DB instances by using snapshots and Amazon S3.

E. Configure an endpoint in AWS Global Accelerator with the two ALBs as equal weighted targets. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group values. Create an Amazon CloudWatch alarm that is based on the HTTPCode_Target_5XX_Count metric for the ALB in the primary Region. Configure the CloudWatch alarm to invoke the Lambda function.

Correct Answer: B
Section:
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.

QUESTION 8
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.)

www.VCEplus.io

A. Use an Elastic Load Balancer to distribute traffic across multiple EC2 instances. Ensure that the EC2 instances are part of an Auto Scaling group that has a minimum capacity of two instances.

B. Use an Elastic Load Balancer to distribute traffic across multiple EC2 instances Ensure that the EC2 instances are configured in unlimited mode.

C. Modify the DB instance to create a read replica in the same Availability Zone. Promote the read replica to be the primary DB instance in failure scenarios.

D. Modify the DB instance to create a Multi-AZ deployment that extends across two Availability Zones.

E. Create a replication group for the ElastiCache for Redis cluster. Configure the cluster to use an Auto Scaling group that has a minimum capacity of two instances.

F. Create a replication group for the ElastiCache for Redis cluster. Enable Multi-AZ on the cluster.

Correct Answer: A, D, F
Section:
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 failover to one of the replicas if the primary node fails or becomes unreachable6

QUESTION 9
A retail company is operating its ecommerce application on AWS. The application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The company uses an Amazon RDS DB instance as the database backend. Amazon CloudFront is configured with one origin that points to the ALB. Static content is cached. Amazon Route 53 is used to host all public zones.

CEplus

After an update of the application, the ALB occasionally returns a 502 status code (Bad Gateway) error. The root cause is malformed HTTP headers that are returned to the ALB. The webpage returns successfully when a solutions architect reloads the webpage immediately after the error occurs.
While the company is working on the problem, the solutions architect needs to provide a custom error page instead of the standard ALB error page to visitors.
Which combination of steps will meet this requirement with the LEAST amount of operational overhead? (Choose two.)

A. Create an Amazon S3 bucket. Configure the S3 bucket to host a static webpage. Upload the custom error pages to Amazon S3.
B. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function if the ALB health check response Target.FailedHealthChecks is greater than 0. Configure the Lambda function to modify the forwarding rule at the ALB to point to a publicly accessible web server.
C. Modify the existing Amazon Route 53 records by adding health checks. Configure a fallback target if the health check fails. Modify DNS records to point to a publicly accessible webpage.
D. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function if the ALB health check response Elb.InternalError is greater than 0. Configure the Lambda function to modify the forwarding rule at the ALB to point to a public accessible web server.
E. Add a custom error response by configuring a CloudFront custom error page. Modify DNS records to point to a publicly accessible web page.

Correct Answer: C, E
Section:
Explanation:
'Save your custom error pages in a location that is accessible to CloudFront. We recommend that you store them in an Amazon S3 bucket, and that you don't store them in the same place as the rest of your website or application's content. If you store the custom error pages on the same origin as your website or application, and the origin starts to return 5xx errors, CloudFront can't get the custom error pages because the origin server is unavailable.' https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/GeneratingCustomErrorResponses.html

QUESTION 10
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 Organizations. Configure the VPCs to share the same CIDR range and subnets as the VPC in the infrastructure account. Peer the VPCs in each individual account with the VPC in the infrastructure account,
D. Create a resource share in AWS Resource Access Manager in the infrastructure account. Select the specific AWS Organizations OU that will use the shared network. Select each subnet to associate with the resource share.
E. Create a resource share in AWS Resource Access Manager in the infrastructure account. Select the specific AWS Organizations OU that will use the shared network. Select each prefix list to associate with the resource share.

Correct Answer: A, E
Section:
Explanation:
https://docs.aws.amazon.com/vpc/latest/userguide/sharing-managed-prefix-lists.html

QUESTION 11
A company wants to use a third-party software-as-a-service (SaaS) application. The third-party SaaS application is consumed through several API calls. The third-party SaaS application also runs on AWS inside a VPC.
The company will consume the third-party SaaS application from inside a VPC. The company has internal security policies that mandate the use of private connectivity that does not traverse the internet. No resources that run in the company VPC are allowed to be accessed from outside the company's VPC. All permissions must conform to the principles of least privilege.

VCEûp CEplus IT Certification Exams - Questions & Answers | VCEplus.io CEplus VCEûp

CEplus

Which solution meets these requirements?

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

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

QUESTION 12
A company needs to implement a patching process for its servers. The on-premises servers and Amazon EC2 instances use a variety of tools to perform patching. Management requires a single report showing the patch status of all the servers and instances.
Which set of actions should a solutions architect take to meet these requirements?

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

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

QUESTION 13
A company is running an application on several Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer. The load on the application varies throughout the day, and EC2 instances are scaled in and out on a regular basis. Log files from the EC2 instances are copied to a central Amazon S3 bucket every 15 minutes. The security team discovers that log files are missing from some of the terminated EC2 instances.
Which set of actions will ensure that log files are copied to the central S3 bucket from the terminated EC2 instances?

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

VCEplus

topic. From the SNS notification, call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send ABANDON to the Auto Scaling group to terminate the instance.

Correct Answer: B
Section:
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-instance/
https://github.com/aws-samples/aws-lambda-lifecycle-hooks-function
https://github.com/aws-samples/aws-lambda-lifecycle-hooks-function/blob/master/cloudformation/template.yaml

QUESTION 14
A company is using multiple AWS accounts The DNS records are stored in a private hosted zone for Amazon Route 53 in Account A The company's applications and databases are running in Account B.
A solutions architect win deploy a two-net application In a new VPC To simplify the configuration, the db.example com CNAME record set tor the Amazon RDS endpoint was created in a private hosted zone for Amazon Route 53.
During deployment, the application failed to start. Troubleshooting revealed that db.example com is not resolvable on the Amazon EC2 instance The solutions architect confirmed that the record set was created correctly in Route 53.
Which combination of steps should the solutions architect take to resolve this issue? (Select TWO)

A. Deploy the database on a separate EC2 instance in the new VPC Create a record set for the instance's private IP in the private hosted zone
B. Use SSH to connect to the application tier EC2 instance Add an RDS endpoint IP address to the /eto/resolv.conf file
C. Create an authorization lo associate the private hosted zone in Account A with the new VPC In Account B
D. Create a private hosted zone for the example.com domain m Account B Configure Route 53 replication between AWS accounts
E. Associate a new VPC in Account B with a hosted zone in Account A. Delete the association authorization In Account A.

Correct Answer: C, E
Section:
Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/private-hosted-zone-different-account/

QUESTION 15
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 storage. Copy the site contents to the volumes at launch and to Amazon S3 al shutdown.
C. Configure an Amazon CloudFront distribution. Point the distribution to an S3 bucket, and migrate the videos from EFS to Amazon S3.
D. Set up an Amazon CloudFront distribution for all site contents, and point the distribution at the ALB.

Correct Answer: C
Section:
Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-https-connection-fails/

VCEûp VCEplus IT Certification Exams - Questions & Answers | VCEplus.io VCEplus VCEûp

VCEplus

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

QUESTION 16
A company with global offices has a single 1 Gbps AWS Direct Connect connection to a single AWS Region. The company's on-premises network uses the connection to communicate with the company's resources in the AWS Cloud. The connection has a single private virtual interface that connects to a single VPC.
A solutions architect must implement a solution that adds a redundant Direct Connect connection in the same Region. The solution also must provide connectivity to other Regions through the same pair of Direct Connect connections as the company expands into other Regions.
Which solution meets these requirements?

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

Correct Answer: A
Section:
Explanation:
A Direct Connect gateway is a globally available resource. You can create the Direct Connect gateway in any Region and access it from all other Regions. The following describe scenarios where you can use a Direct Connect gateway. https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-gateways-intro.html

QUESTION 17
A company is developing a new service that will be accessed using TCP on a static port A solutions architect must ensure that the service is highly available, has redundancy across Availability Zones, and is accessible using the DNS name myservice.com, which is publicly accessible The service must use fixed address assignments so other companies can add the addresses to their allow lists.
Assuming that resources are deployed in multiple Availability Zones in a single Region, which solution will meet these requirements?

A. Create Amazon EC2 instances with an Elastic IP address for each instance Create a Network Load Balancer (NLB) and expose the static TCP port Register EC2 instances with the NLB Create a new name server record set named my service com, and assign the Elastic IP addresses of the EC2 instances to the record set Provide the Elastic IP addresses of the EC2 instances to the other companies to add to their allow lists
B. Create an Amazon ECS cluster and a service definition for the application Create and assign public IP addresses for the ECS cluster Create a Network Load Balancer (NLB) and expose the TCP port Create a target group and assign the ECS cluster name to the NLB Create a new A record set named my service com and assign the public IP addresses of the ECS cluster to the record set Provide the public IP addresses of the ECS cluster to the other companies to add to their allow lists
C. Create Amazon EC2 instances for the service Create one Elastic IP address for each Availability Zone Create a Network Load Balancer (NLB) and expose the assigned TCP port Assign the Elastic IP addresses to the NLB for each Availability Zone Create a target group and register the EC2 instances with the NLB Create a new A (alias) record set named my service com, and assign the NLB DNS name to the record set.
D. Create an Amazon ECS cluster and a service definition for the application Create and assign public IP address for each host in the cluster Create an Application Load Balancer (ALB) and expose the static TCP port Create a target group and assign the ECS service definition name to the ALB Create a new CNAME record set and associate the public IP addresses to the record set Provide the Elastic IP addresses of the Amazon EC2 instances to the other companies to add to their allow lists

Correct Answer: C

VCEûp VCEplus IT Certification Exams - Questions & Answers | VCEplus.io VCEplus VCEûp

VCEplus

Section:
Explanation:
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-to-elb-load-balancer.html
Create a Network Load Balancer (NLB) and expose the assigned TCP port. Assign the Elastic IP addresses to the NLB for each Availability Zone. Create a target group and register the EC2 instances with the NLB. Create a new A (alias) record set named my.service.com, and assign the NLB DNS name to the record set. As it uses the NLB as the resource in the A-record, traffic will be routed through the NLB, and it will automatically route the traffic to the healthy instances based on the health checks and also it provides the fixed address assignments as the other companies can add the NLB's Elastic IP addresses to their allow lists.

QUESTION 18
A company uses an on-premises data analytics platform. The system is highly available in a fully redundant configuration across 12 servers in the company's data center.
The system runs scheduled jobs, both hourly and daily, in addition to one-time requests from users. Scheduled jobs can take between 20 minutes and 2 hours to finish running and have tight SLAs. The scheduled jobs account for 65% of the system usage. User jobs typically finish running in less than 5 minutes and have no SLA. The user jobs account for 35% of system usage. During system failures, scheduled jobs must continue to meet SLAs. However, user jobs can be delayed. A solutions architect needs to move the system to Amazon EC2 instances and adopt a consumption-based model to reduce costs with no long-term commitments. The solution must maintain high availability and must not affect the SLAs. Which solution will meet these requirements MOST cost-effectively?

A. Split the 12 instances across two Availability Zones in the chosen AWS Region. Run two instances in each Availability Zone as On-Demand Instances with Capacity Reservations. Run four instances in each Availability Zone as Spot Instances.
B. Split the 12 instances across three Availability Zones in the chosen AWS Region. In one of the Availability Zones, run all four instances as On-Demand Instances with Capacity Reservations. Run the remaining instances as Spot Instances.
C. Split the 12 instances across three Availability Zones in the chosen AWS Region. Run two instances in each Availability Zone as On-Demand Instances with a Savings Plan. Run two instances in each Availability Zone as Spot Instances.
D. Split the 12 instances across three Availability Zones in the chosen AWS Region. Run three instances in each Availability Zone as On-Demand Instances with Capacity Reservations. Run one instance in each Availability Zone as a Spot Instance.

www.VCEplus.io

Correct Answer: D
Section:
Explanation:
By splitting the 12 instances across three Availability Zones, the system can maintain high availability and availability of resources in case of a failure. Option D also uses a combination of On-Demand Instances with Capacity Reservations and Spot Instances, which allows for scheduled jobs to be run on the On-Demand instances with guaranteed capacity, while also taking advantage of the cost savings from Spot Instances for the user jobs which have lower SLA requirements.

QUESTION 19
A security engineer determined that an existing application retrieves credentials to an Amazon RDS for MySQL database from an encrypted file in Amazon S3. For the next version of the application, the security engineer wants to implement the following application design changes to improve security:
The database must use strong, randomly generated passwords stored in a secure AWS managed service.
The application resources must be deployed through AWS CloudFormation.
The application must rotate credentials for the database every 90 days.
A solutions architect will generate a CloudFormation template to deploy the application.
Which resources specified in the CloudFormation template will meet the security engineer's requirements with the LEAST amount of operational overhead?

A. Generate the database password as a secret resource using AWS Secrets Manager. Create an AWS Lambda function resource to rotate the database password. Specify a Secrets Manager RotationSchedule resource to rotate the database password every 90 days.
B. Generate the database password as a SecureString parameter type using AWS Systems Manager Parameter Store. Create an AWS Lambda function resource to rotate the database password. Specify a Parameter Store RotationSchedule resource to rotate the database password every 90 days.
C. Generate the database password as a secret resource using AWS Secrets Manager. Create an AWS Lambda function resource to rotate the database password. Create an Amazon EventBridge scheduled rule resource to trigger the Lambda function password rotation every 90 days.

VCEplus

D. Generate the database password as a SecureString parameter type using AWS Systems Manager Parameter Store. Specify an AWS AppSync DataSource resource to automatically rotate the database password every 90 days.

Correct Answer: B
Section:
Explanation:
https://aws.amazon.com/blogs/security/how-to-securely-provide-database-credentials-to-lambda-functions-by-using-aws-secrets-manager/
https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotating-secrets.html
https://docs.aws.amazon.com/secretsmanager/latest/userguide/integrating_cloudformation.html

QUESTION 20
A company is storing data in several Amazon DynamoDB tables. A solutions architect must use a serverless architecture to make the data accessible publicly through a simple API over HTTPS. The solution must scale automatically in response to demand.
Which solutions meet these requirements? (Choose two.)

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

Correct Answer: A, C
Section:
Explanation:
https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-overview-developer-experience.html

QUESTION 21
A company has registered 10 new domain names. The company uses the domains for online marketing. The company needs a solution that will redirect online visitors to a specific URL for each domain. All domains and target URLs are defined in a JSON document. All DNS records are managed by Amazon Route 53.
A solutions architect must implement a redirect service that accepts HTTP and HTTPS requests.
Which combination of steps should the solutions architect take to meet these requirements with the LEAST amount of operational effort? (Choose three.)

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

Correct Answer: C, E, F
Section:
Explanation:
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-edge-how-it-works-tutorial.html

QUESTION 22
A company that has multiple AWS accounts is using AWS Organizations. The company's AWS accounts host VPCs, Amazon EC2 instances, and containers.

VCEûp CEplus IT Certification Exams - Questions & Answers | VCEplus.io CEplus VCEûp

VCEplus

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

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

QUESTION 23
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 account. Associate the attachment with the transit gateway in the management account by using the transit gateway ID.
D. Launch an AWS CloudFormation stack set from the management account that automatical^ creates a new VPC and a peering transit gateway attachment in a member account. Share the attachment with the transit gateway in the management account by using a transit gateway service-linked role.
E. From the management account, share the transit gateway with member accounts by using AWS Service Catalog

Correct Answer: A, C
Section:
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-transitgatewayattachment.html

QUESTION 24
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.

VCEplus

B. Create an IAM role named procurement-manager-role in all AWS accounts in the organization Add the AdministratorAccess managed policy to the role Define a permissions boundary with the AWSPrivateMarketplaceAdminFullAccess managed policy and attach it to all the developer roles.

C. Create an IAM role named procurement-manager-role in all the shared services accounts in the organization Add the AWSPrivateMarketplaceAdminFullAccess managed policy to the role Create an organization root-level SCP to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role Create another organization root-level SCP to deny permissions to create an IAM role named procurement-manager-role to everyone in the organization.

D. Create an IAM role named procurement-manager-role in all AWS accounts that will be used by developers. Add the AWSPrivateMarketplaceAdminFullAccess managed policy to the role. Create an SCP in Organizations to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role. Apply the SCP to all the shared services accounts in the organization.

Correct Answer: C
Section:
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-using-iam-and-aws-organizations/
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.

QUESTION 25
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

VCEplus

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

Correct Answer: B
Section:
Explanation:
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_inheritance_auth.html

QUESTION 26
A company is hosting a monolithic REST-based API for a mobile app on five Amazon EC2 instances in public subnets of a VPC. Mobile clients connect to the API by using a domain name that is hosted on Amazon Route 53. The company has created a Route 53 multivalue answer routing policy with the IP addresses of all the EC2 instances. Recently, the app has been overwhelmed by large and sudden increases to traffic. The app has not been able to keep up with the traffic.
A solutions architect needs to implement a solution so that the app can handle the new and varying load.
Which solution will meet these requirements with the LEAST operational overhead?

A. Separate the API into individual AWS Lambda functions. Configure an Amazon API Gateway REST API with Lambda integration for the backend. Update the Route 53 record to point to the API Gateway API.
B. Containerize the API logic. Create an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. Run the containers in the cluster by using Amazon EC2. Create a Kubernetes ingress. Update the Route 53 record to point to the Kubernetes ingress.
C. Create an Auto Scaling group. Place all the EC2 instances in the Auto Scaling group. Configure the Auto Scaling group to perform scaling actions that are based on CPU utilization. Create an AWS Lambda function that reacts to Auto Scaling group changes and updates the Route 53 record.
D. Create an Application Load Balancer (ALB) in front of the API. Move the EC2 instances to private subnets in the VPC. Add the EC2 instances as targets for the ALB. Update the Route 53 record to point to the ALB.

Correct Answer: D
Section:
Explanation:
By breaking down the monolithic API into individual Lambda functions and using API Gateway to handle the incoming requests, the solution can automatically scale to handle the new and varying load without the need for manual scaling actions. Additionally, this option will automatically handle the traffic without the need of having EC2 instances running all the time and only pay for the number of requests and the duration of the execution of the Lambda function.
By updating the Route 53 record to point to the API Gateway, the solution can handle the traffic and also it will direct the traffic to the correct endpoint.

QUESTION 27
A company has created an OU in AWS Organizations for each of its engineering teams Each OU owns multiple AWS accounts. The organization has hundreds of AWS accounts A solutions architect must design a solution so that each OU can view a breakdown of usage costs across its AWS accounts. Which solution meets these requirements?

A. Create an AWS Cost and Usage Report (CUR) for each OU by using AWS Resource Access Manager Allow each team to visualize the CUR through an Amazon QuickSight dashboard.
B. Create an AWS Cost and Usage Report (CUR) from the AWS Organizations management account- Allow each team to visualize the CUR through an Amazon QuickSight dashboard
C. Create an AWS Cost and Usage Report (CUR) in each AWS Organizations member account Allow each team to visualize the CUR through an Amazon QuickSight dashboard.
D. Create an AWS Cost and Usage Report (CUR) by using AWS Systems Manager Allow each team to visualize the CUR through Systems Manager OpsCenter dashboards

Correct Answer: B
Section:
Explanation:
https://docs.aws.amazon.com/cur/latest/userguide/billing-cur-limits.html

QUESTION 28

VCEûp VCEplus IT Certification Exams - Questions & Answers | VCEplus.io VCEplus VCEûp

VCEplus

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

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

Correct Answer: B
Section:
Explanation:
https://aws.amazon.com/storagegateway/file/
https://docs.aws.amazon.com/fsx/latest/WindowsGuide/migrate-files-to-fsx-datasync.html
https://docs.aws.amazon.com/systems-manager/latest/userguide/prereqs-operating-systems.html#prereqs-os-windows-server

QUESTION 29
A company's solutions architect is reviewing a web application that runs on AWS. The application references static assets in an Amazon S3 bucket in the us-east-1 Region. The company needs resiliency across multiple AWS Regions. The company already has created an S3 bucket in a second Region.
Which solution will meet these requirements with the LEAST operational overhead?

A. Configure the application to write each object to both S3 buckets. Set up an Amazon Route 53 public hosted zone with a record set by using a weighted routing policy for each S3 bucket. Configure the application to reference the objects by using the Route 53 DNS name.
B. Create an AWS Lambda function to copy objects from the S3 bucket in us-east-1 to the S3 bucket in the second Region. Invoke the Lambda function each time an object is written to the S3 bucket in us-east-1. Set up an Amazon CloudFront distribution with an origin group that contains the two S3 buckets as origins.
C. Configure replication on the S3 bucket in us-east-1 to replicate objects to the S3 bucket in the second Region Set up an Amazon CloudFront distribution with an origin group that contains the two S3 buckets as origins.
D. Configure replication on the S3 bucket in us-east-1 to replicate objects to the S3 bucket in the second Region. If failover is required, update the application code to load S3 objects from the S3 bucket in the second Region.

Correct Answer: C
Section:
Explanation:
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/high_availability_origin_failover.html

QUESTION 30
A company is hosting a three-tier web application in an on-premises environment. Due to a recent surge in traffic that resulted in downtime and a significant financial impact, company management has ordered that the application be moved to AWS. The application is written in .NET and has a dependency on a MySQL database A solutions architect must design a scalable and highly available solution to meet the demand of 200000 daily users.
Which steps should the solutions architect take to design an appropriate solution?

A. Use AWS Elastic Beanstalk to create a new application with a web server environment and an Amazon RDS MySQL Multi-AZ DB instance The environment should launch a Network Load Balancer (NLB) in front of an Amazon EC2 Auto Scaling group in multiple Availability Zones Use an Amazon Route 53 alias record to route traffic from the company's domain to the NLB.
B. Use AWS CloudFormation to launch a stack containing an Application Load Balancer (ALB) in front of an Amazon EC2 Auto Scaling group spanning three Availability Zones. The stack should launch a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a Retain deletion policy. Use an Amazon Route 53 alias record to route traffic from the company's domain to the ALB

VCEûp VCEplus IT Certification Exams - Questions & Answers | VCEplus.io VCEplus VCEûp

VCEplus

C. Use AWS Elastic Beanstalk to create an automatically scaling web server environment that spans two separate Regions with an Application Load Balancer (ALB) in each Region. Create a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a cross-Region read replica Use Amazon Route 53 with a geoproximity routing policy to route traffic between the two Regions.
D. Use AWS CloudFormation to launch a stack containing an Application Load Balancer (ALB) in front of an Amazon ECS cluster of Spot Instances spanning three Availability Zones The stack should launch an Amazon RDS MySQL DB instance with a Snapshot deletion policy Use an Amazon Route 53 alias record to route traffic from the company's domain to the ALB

Correct Answer: C
Section:
Explanation:
Using AWS CloudFormation to launch a stack with an Application Load Balancer (ALB) in front of an Amazon EC2 Auto Scaling group spanning three Availability Zones, a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a Retain deletion policy, and an Amazon Route 53 alias record to route traffic from the company's domain to the ALB will ensure that

QUESTION 31
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 accounts. Use service-managed permissions. Set deployment options to deploy to an organization. Use CloudFormation StackSets drift detection.
B. Create stacks in the Organizations member accounts. Use self-service permissions. Set deployment options to deploy to an organization. Enable the CloudFormation StackSets automatic deployment.
C. Create a stack set in the Organizations management account Use service-managed permissions. Set deployment options to deploy to the organization. Enable CloudFormation StackSets automatic deployment.
D. Create stacks in the Organizations management account. Use service-managed permissions. Set deployment options to deploy to the organization. Enable CloudFormation StackSets drift detection.

Correct Answer: C
Section:
Explanation:
https://aws.amazon.com/blogs/aws/use-cloudformation-stacksets-to-provision-resources-across-multiple-aws-accounts-and-regions/

QUESTION 32
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 Advisor. Follow the recommendations for cost optimization.

Correct Answer: A, D, E
Section:
Explanation:
https://docs.aws.amazon.com/application-discovery/latest/userguide/discovery-agent.html https://docs.aws.amazon.com/migrationhub/latest/ug/ec2-recommendations.html

VCEûp VCEplus IT Certification Exams - Questions & Answers | VCEplus.io VCEplus VCEûp

VCEplus

QUESTION 33
A company is hosting an image-processing service on AWS in a VPC. The VPC extends across two Availability Zones. Each Availability Zone contains one public subnet and one private subnet.
The service runs on Amazon EC2 instances in the private subnets. An Application Load Balancer in the public subnets is in front of the service. The service needs to communicate with the internet and does so through two NAT gateways. The service uses Amazon S3 for image storage. The EC2 instances retrieve approximately 1 ' of data from an S3 bucket each day.
The company has promoted the service as highly secure. A solutions architect must reduce cloud expenditures as much as possible without compromising the service's security posture or increasing the time spent on ongoing operations.
Which solution will meet these requirements?

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

Correct Answer: C
Section:
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.

QUESTION 34
A company recently deployed an application on AWS. The application uses Amazon DynamoDB. The company measured the application load and configured the RCUs and WCUs on the DynamoDB table to match the expected peak load. The peak load occurs once a week for a 4-hour period and is double the average load. The application load is close to the average load tor the rest of the week. The access pattern includes many more writes to the table than reads of the table.
A solutions architect needs to implement a solution to minimize the cost of the table.
Which solution will meet these requirements?

www.VCEplus.io

A. Use AWS Application Auto Scaling to increase capacity during the peak period. Purchase reserved RCUs and WCUs to match the average load.
B. Configure on-demand capacity mode for the table.
C. Configure DynamoDB Accelerator (DAX) in front of the table. Reduce the provisioned read capacity to match the new peak load on the table.
D. Configure DynamoDB Accelerator (DAX) in front of the table. Configure on-demand capacity mode for the table.

Correct Answer: D
Section:
Explanation:
This solution meets the requirements by using Application Auto Scaling to automatically increase capacity during the peak period, which will handle the double the average load. And by purchasing reserved RCUs and WCUs to match the average load, it will minimize the cost of the table for the rest of the week when the load is close to the average.

QUESTION 35
A solutions architect needs to advise a company on how to migrate its on-premises data processing application to the AWS Cloud. Currently, users upload input files through a web portal. The web server then stores the uploaded files on NAS and messages the processing server over a message queue. Each media file can take up to 1 hour to process. The company has determined that the number of media files awaiting processing is significantly higher during business hours, with the number of files rapidly declining after business hours.
What is the MOST cost-effective migration recommendation?

A. Create a queue using Amazon SQS. Configure the existing web server to publish to the new queue. When there are messages in the queue, invoke an AWS Lambda function to pull requests from the queue and process the files. Store the processed files in an Amazon S3 bucket.
B. Create a queue using Amazon M. Configure the existing web server to publish to the new queue. When there are messages in the queue, create a new Amazon EC2 instance to pull requests from the queue and process the files. Store the processed files in Amazon EFS. Shut down the EC2 instance after the task is complete.

CEplus

C. Create a queue using Amazon MO. Configure the existing web server to publish to the new queue. When there are messages in the queue, invoke an AWS Lambda function to pull requests from the queue and process the files. Store the processed files in Amazon EFS.

D. Create a queue using Amazon SOS. Configure the existing web server to publish to the new queue. Use Amazon EC2 instances in an EC2 Auto Scaling group to pull requests from the queue and process the files. Scale the EC2 instances based on the SOS queue length. Store the processed files in an Amazon S3 bucket.

Correct Answer: D
Section:
Explanation:
https://aws.amazon.com/blogs/compute/operating-lambda-performance-optimization-part-1/

QUESTION 36
A company is using Amazon OpenSearch Service to analyze data. The company loads data into an OpenSearch Service cluster with 10 data nodes from an Amazon S3 bucket that uses S3 Standard storage. The data resides in the cluster for 1month for read-only analysis. After 1month, the company deletes the index that contains the data from the cluster. For compliance purposes, the company must retain a copy of all input data.
The company is concerned about ongoing costs and asks a solutions architect to recommend a new solution.
Which solution will meet these requirements MOST cost-effectively?

A. Replace all the data nodes with UltraWarm nodes to handle the expected capacity. Transition the input data from S3 Standard to S3 Glacier Deep Archive when the company loads the data into the cluster.

B. Reduce the number of data nodes in the cluster to 2 Add UltraWarm nodes to handle the expected capacity. Configure the indexes to transition to UltraWarm when OpenSearch Service ingests the data. Transition the input data to S3 Glacier Deep Archive after 1month by using an S3 Lifecycle policy.

C. Reduce the number of data nodes in the cluster to 2. Add UltraWarm nodes to handle the expected capacity. Configure the indexes to transition to UltraWarm when OpenSearch Service ingests the data. Add cold storage nodes to the cluster Transition the indexes from UltraWarm to cold storage. Delete the input data from the S3 bucket after 1month by using an S3 Lifecycle policy.

D. Reduce the number of data nodes in the cluster to 2. Add instance-backed data nodes to handle the expected capacity. Transition the input data from S3 Standard to S3 Glacier Deep Archive when the company loads the data into the cluster.

Correct Answer: B
Section:
Explanation:
By reducing the number of data nodes in the cluster to 2 and adding UltraWarm nodes to handle the expected capacity, the company can reduce the cost of running the cluster. Additionally, configuring the indexes to transition to UltraWarm when OpenSearch Service ingests the data will ensure that the data is stored in the most cost-effective manner. Finally, transitioning the input data to S3 Glacier Deep Archive after 1 month by using an S3 Lifecycle policy will ensure that the data is retained for compliance purposes, while also reducing the ongoing costs.

QUESTION 37
A company has 10 accounts that are part of an organization in AWS Organizations AWS Config is configured in each account All accounts belong to either the Prod OU or the NonProd OU
The company has set up an Amazon EventBridge rule in each AWS account to notify an Amazon Simple Notification Service (Amazon SNS) topic when an Amazon EC2 security group inbound rule is created with 0.0.0.0/0 as the source The company's security team is subscribed to the SNS topic
For all accounts in the NonProd OU the security team needs to remove the ability to create a security group inbound rule that includes 0.0.0.0/0 as the source
Which solution will meet this requirement with the LEAST operational overhead?

A. Modify the EventBridge rule to invoke an AWS Lambda function to remove the security group inbound rule and to publish to the SNS topic Deploy the updated rule to the NonProd OU

B. Add the vpc-sg-open-only-to-authorized-ports AWS Config managed rule to the NonProd OU

C. Configure an SCP to allow the ec2 AulhonzeSecurityGroupIngress action when the value of the aws SourceIp condition key is not 0.0.0.0/0 Apply the SCP to the NonProd OU

D. Configure an SCP to deny the ec2 AuthorizeSecurityGroupIngress action when the value of the aws SourceIp condition key is 0.0.0.0/0 Apply the SCP to the NonProd OU

VCEup CEplus IT Certification Exams - Questions & Answers | VCEplus.io CEplus VCEup

VCEplus

Correct Answer: D
Section:
Explanation:
This solution will meet the requirement with the least operational overhead because it directly denies the creation of the security group inbound rule with 0.0.0.0/0 as the source, which is the exact requirement. Additionally, it does not require any additional steps or resources such as invoking a Lambda function or adding a Config rule.
An SCP (Service Control Policy) is a policy that you can use to set fine-grained permissions for your AWS accounts within your organization. You can use SCPs to set permissions for the root user of an account and to delegate permissions to IAM users and roles in the accounts. You can use SCPs to set permissions that allow or deny access to specific services, actions, and resources.
To implement this solution, you would need to create an SCP that denies the ec2:AuthorizeSecurityGroupIngress action when the value of the aws:SourceIp condition key is 0.0.0.0/0. This SCP would then be applied to the NonProd OU. This would ensure that any security group inbound rule that includes 0.0.0.0/0 as the source will be denied, thus meeting the requirement.
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_condition-keys.html

QUESTION 38
A company hosts a Git repository in an on-premises data center. The company uses webhooks to invoke functionality that runs in the AWS Cloud. The company hosts the webhook logic on a set of Amazon EC2 instances in an Auto Scaling group that the company set as a target for an Application Load Balancer (ALB). The Git server calls the ALB for the configured webhooks. The company wants to move the solution to a serverless architecture.
Which solution will meet these requirements with the LEAST operational overhead?

A. For each webhook, create and configure an AWS Lambda function URL. Update the Git servers to call the individual Lambda function URLs.
B. Create an Amazon API Gateway HTTP API. Implement each webhook logic in a separate AWS Lambda function. Update the Git servers to call the API Gateway endpoint.
C. Deploy the webhook logic to AWS App Runner. Create an ALB, and set App Runner as the target. Update the Git servers to call the ALB endpoint.
D. Containerize the webhook logic. Create an Amazon Elastic Container Service (Amazon ECS) cluster, and run the webhook logic in AWS Fargate. Create an Amazon API Gateway REST API, and set Fargate as the target. Update the Git servers to call the API Gateway endpoint.

www.VCEplus.io

Correct Answer: B
Section:
Explanation:
https://aws.amazon.com/solutions/implementations/git-to-s3-using-webhooks/
https://medium.com/mindorks/building-webhook-is-easy-using-aws-lambda-and-api-gateway-56f5e5c3a596

QUESTION 39
A company is planning to migrate 1,000 on-premises servers to AWS. The servers run on several VMware clusters in the company's data center. As part of the migration plan, the company wants to gather server metrics such as CPU details, RAM usage, operating system information, and running processes. The company then wants to query and analyze the data.
Which solution will meet these requirements?

A. Deploy and configure the AWS Agentless Discovery Connector virtual appliance on the on-premises hosts. Configure Data Exploration in AWS Migration Hub. Use AWS Glue to perform an ETL job against the data. Query the data by using Amazon S3 Select.
B. Export only the VM performance information from the on-premises hosts. Directly import the required data into AWS Migration Hub. Update any missing information in Migration Hub. Query the data by using Amazon QuickSight.
C. Create a script to automatically gather the server information from the on-premises hosts. Use the AWS CLI to run the put-resource-attributes command to store the detailed server data in AWS Migration Hub. Query the data directly in the Migration Hub console.
D. Deploy the AWS Application Discovery Agent to each on-premises server. Configure Data Exploration in AWS Migration Hub. Use Amazon Athena to run predefined queries against the data in Amazon S3.

Correct Answer: D
Section:
Explanation:
it covers all the requirements mentioned in the question, it will allow collecting the detailed metrics, including process information and it provides a way to query and analyze the data using Amazon Athena.

VCEplus

QUESTION 40
A company is building a serverless application that runs on an AWS Lambda function that is attached to a VPC. The company needs to integrate the application with a new service from an external provider. The external provider supports only requests that come from public IPv4 addresses that are in an allow list.
The company must provide a single public IP address to the external provider before the application can start using the new service.
Which solution will give the application the ability to access the new service?

A. Deploy a NAT gateway. Associate an Elastic IP address with the NAT gateway. Configure the VPC to use the NAT gateway.
B. Deploy an egress-only internet gateway. Associate an Elastic IP address with the egress-only internet gateway. Configure the elastic network interface on the Lambda function to use the egress-only internet gateway.
C. Deploy an internet gateway. Associate an Elastic IP address with the internet gateway. Configure the Lambda function to use the internet gateway.
D. Deploy an internet gateway. Associate an Elastic IP address with the internet gateway. Configure the default route in the public VPC route table to use the internet gateway.

Correct Answer: A
Section:
Explanation:
This solution will give the Lambda function access to the internet by routing its outbound traffic through the NAT gateway, which has a public Elastic IP address. This will allow the external provider to whitelist the single public IP address associated with the NAT gateway, and enable the application to access the new service
Deploying a NAT gateway and associating an Elastic IP address with it, and then configuring the VPC to use the NAT gateway, will give the application the ability to access the new service. This is because the NAT gateway will be the single public IP address that the external provider needs for the allow list. The NAT gateway will allow the application to access the service, while keeping the underlying Lambda functions private.
When configuring NAT gateways, you should ensure that the route table associated with the NAT gateway has a route to the internet gateway with a target of the internet gateway. Additionally, you should ensure that the security group associated with the NAT gateway allows outbound traffic from the Lambda functions.
AWS Certified Solutions Architect Professional Official Amazon Text Book [1], page 456
https://docs.aws.amazon.com/vpc/latest/userguide/VPC_NAT_Gateway.html

www.VCEplus.io

QUESTION 41
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.

Correct Answer: B, D
Section:
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

QUESTION 42

VCEûp VCEplus IT Certification Exams - Questions & Answers | VCEplus.io VCEplus VCEûp

CEplus

A company is planning to host a web application on AWS and works to load balance the traffic across a group of Amazon EC2 instances. One of the security requirements is to enable end-to-end encryption in transit between the client and the web server.
Which solution will meet this requirement?

A. Place the EC2 instances behind an Application Load Balancer (ALB) Provision an SSL certificate using AWS Certificate Manager (ACM), and associate the SSL certificate with the ALB. Export the SSL certificate and install it on each EC2 instance. Configure the ALB to listen on port 443 and to forward traffic to port 443 on the instances.
B. Associate the EC2 instances with a target group. Provision an SSL certificate using AWS Certificate Manager (ACM). Create an Amazon CloudFront distribution and configure It to use the SSL certificate. Set CloudFront to use the target group as the origin server
C. Place the EC2 instances behind an Application Load Balancer (ALB). Provision an SSL certificate using AWS Certificate Manager (ACM), and associate the SSL certificate with the ALB. Provision a third-party SSL certificate and install it on each EC2 instance. Configure the ALB to listen on port 443 and to forward traffic to port 443 on the instances.
D. Place the EC2 instances behind a Network Load Balancer (NLB). Provision a third-party SSL certificate and install it on the NLB and on each EC2 instance. Configure the NLB to listen on port 443 and to forward traffic to port 443 on the instances.

Correct Answer: A
Section:
Explanation:
Option A is correct because placing the EC2 instances behind an Application Load Balancer (ALB) and associating an SSL certificate from AWS Certificate Manager (ACM) with the ALB enables encryption in transit between the client and the ALB. Exporting the SSL certificate and installing it on each EC2 instance enables encryption in transit between the ALB and the web server. Configuring the ALB to listen on port 443 and to forward traffic to port 443 on the instances ensures that HTTPS is used for both connections.This solution achieves end-to-end encryption in transit for the web application12

QUESTION 43
A company wants to migrate its data analytics environment from on premises to AWS The environment consists of two simple Node js applications One of the applications collects sensor data and loads it into a MySQL database The other application aggregates the data into reports When the aggregation jobs run. some of the load jobs fail to run correctly
The company must resolve the data loading issue The company also needs the migration to occur without interruptions or changes for the company's customers
What should a solutions architect do to meet these requirements?

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

Correct Answer: C
Section:
Explanation:
Set up an Amazon Aurora MySQL database. Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora. Create an Aurora Replica for the Aurora MySQL database, and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as AWS Lambda functions behind an Application Load Balancer (ALB), and use Amazon RDS Proxy to write to the Aurora MySQL database. When the databases are synced, point the collector DNS record to the ALB. Disable the AWS DMS sync task after the cutover from on premises to AWS.
Amazon RDS Proxy allows applications to pool and share connections established with the database, improving database efficiency and application scalability. With RDS Proxy, failover times for Aurora and RDS databases are reduced by up to 66%

VCEûp CEplus IT Certification Exams - Questions & Answers | VCEplus.io CEplus VCEûp

CEplus

QUESTION 44
A health insurance company stores personally identifiable information (PII) in an Amazon S3 bucket. The company uses server-side encryption with S3 managed encryption keys (SSE-S3) to encrypt the objects. According to a new requirement, all current and future objects in the S3 bucket must be encrypted by keys that the company's security team manages. The S3 bucket does not have versioning enabled.
Which solution will meet these requirements?

A. In the S3 bucket properties, change the default encryption to SSE-S3 with a customer managed key. Use the AWS CLI to re-upload all objects in the S3 bucket. Set an S3 bucket policy to deny unencrypted PutObject requests.
B. In the S3 bucket properties, change the default encryption to server-side encryption with AWS KMS managed encryption keys (SSE-KMS). Set an S3 bucket policy to deny unencrypted PutObject requests. Use the AWS CLI to re-upload all objects in the S3 bucket.
C. In the S3 bucket properties, change the default encryption to server-side encryption with AWS KMS managed encryption keys (SSE-KMS). Set an S3 bucket policy to automatically encrypt objects on GetObject and PutObject requests.
D. In the S3 bucket properties, change the default encryption to AES-256 with a customer managed key. Attach a policy to deny unencrypted PutObject requests to any entities that access the S3 bucket. Use the AWS CLI to re-upload all objects in the S3 bucket.

Correct Answer: D
Section:
Explanation:
https://docs.aws.amazon.com/AmazonS3/latest/userguide/ServerSideEncryptionCustomerKeys.html Clearly says we need following header for SSE-C x-amz-server-side-encryption-customer-algorithm Use this header to specify the encryption algorithm. The header value must be AES256.

QUESTION 45
A company is running a web application in the AWS Cloud. The application consists of dynamic content that is created on a set of Amazon EC2 instances. The EC2 instances run in an Auto Scaling group that is configured as a target group for an Application Load Balancer (ALB).
The company is using an Amazon CloudFront distribution to distribute the application globally. The CloudFront distribution uses the ALB as an origin. The company uses Amazon Route 53 for DNS and has created an A record of www.example.com for the CloudFront distribution.
A solutions architect must configure the application so that itis highly available and fault tolerant.
Which solution meets these requirements?

A. Provision a full, secondary application deployment in a different AWS Region. Update the Route 53 A record to be a failover record. Add both of the CloudFront distributions as values. Create Route 53 health checks.
B. Provision an ALB, an Auto Scaling group, and EC2 instances in a different AWS Region. Update the CloudFront distribution, and create a second origin for the new ALB. Create an origin group for the two origins. Configure one origin as primary and one origin as secondary.
C. Provision an Auto Scaling group and EC2 instances in a different AWS Region. Create a second target for the new Auto Scaling group in the ALB. Set up the failover routing algorithm on the ALB.
D. Provision a full, secondary application deployment in a different AWS Region. Create a second CloudFront distribution, and add the new application setup as an origin. Create an AWS Global Accelerator accelerator. Add both of the CloudFront distributions as endpoints.

Correct Answer: B
Section:
Explanation:
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistS3AndCustomOrigins.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/high_availability_origin_failover.html
You can set up CloudFront with origin failover for scenarios that require high availability. To get started, you create an origin group with two origins: a primary and a secondary. If the primary origin is unavailable, or returns specific HTTP response status codes that indicate a failure, CloudFront automatically switches to the secondary origin.

QUESTION 46
A company has an organization in AWS Organizations that has a large number of AWS accounts. One of the AWS accounts is designated as a transit account and has a transit gateway that is shared with all of the other AWS accounts AWS Site-to-Site VPN connections are configured between ail of the company's global offices and the transit account The company has AWS Config enabled on all of its accounts.

VCEûp CEplus IT Certification Exams - Questions & Answers | VCEplus.io CEplus VCEûp

CEplus

The company's networking team needs to centrally manage a list of internal IP address ranges that belong to the global offices Developers Will reference this list to gain access to applications securely.
Which solution meets these requirements with the LEAST amount of operational overhead?

A. Create a JSON file that is hosted in Amazon S3 and that lists all of the internal IP address ranges Configure an Amazon Simple Notification Service (Amazon SNS) topic in each of the accounts that can be involved when the JSON file is updated. Subscribe an AWS Lambda function to the SNS topic to update all relevant security group rules with Vie updated IP address ranges.
B. Create a new AWS Config managed rule that contains all of the internal IP address ranges Use the rule to check the security groups in each of the accounts to ensure compliance with the list of IP address ranges. Configure the rule to automatically remediate any noncompliant security group that is detected.
C. In the transit account, create a VPC prefix list with all of the internal IP address ranges. Use AWS Resource Access Manager to share the prefix list with all of the other accounts. Use the shared prefix list to configure security group rules is the other accounts.
D. In the transit account create a security group with all of the internal IP address ranges. Configure the security groups in me other accounts to reference the transit account's security group by using a nested security group reference of *<transit-account-id>./sg-1a2b3c4d'.

Correct Answer: C
Section:
Explanation:
Customer-managed prefix lists --- Sets of IP address ranges that you define and manage. You can share your prefix list with other AWS accounts, enabling those accounts to reference the prefix list in their own resources. https://docs.aws.amazon.com/vpc/latest/userguide/managed-prefix-lists.html
a VPC prefix list is created in the transit account with all of the internal IP address ranges, and then shared to all of the other accounts using AWS Resource Access Manager. This allows for central management of the IP address ranges, and eliminates the need for manual updates to security group rules in each account. This solution also allows for compliance checks to be run using AWS Config and for any non-compliant security groups to be automatically remediated.

QUESTION 47
A company runs a new application as a static website in Amazon S3. The company has deployed the application to a production AWS account and uses Amazon CloudFront to deliver the website. The website calls an Amazon API Gateway REST API. An AWS Lambda function backs each API method.
The company wants to create a CSV report every 2 weeks to show each API Lambda function's recommended configured memory, recommended cost, and the price difference between current configurations and the recommendations. The company will store the reports in an S3 bucket.
Which solution will meet these requirements with the LEAST development time?fs20

A. Create a Lambda function that extracts metrics data for each API Lambda function from Amazon CloudWatch Logs for the 2-week penod_ Collate the data into tabular format. Store the data as a _csvfile in an S3 bucket. Create an Amazon Eventaridge rule to schedule the Lambda function to run every 2 weeks.
B. Opt in to AWS Compute Optimizer. Create a Lambda function that calls the ExportLambdaFunctionRecommendatlons operation. Export the _csv file to an S3 bucket. Create an Amazon Eventaridge rule to schedule the Lambda function to run every 2 weeks.
C. Opt in to AWS Compute Optimizer. Set up enhanced infrastructure metrics. Within the Compute Optimizer console, schedule a job to export the Lambda recommendations to a _csvfile_ Store the file in an S3 bucket every 2 weeks.
D. Purchase the AWS Business Support plan for the production account. Opt in to AWS Compute Optimizer for AWS Trusted Advisor checks. In the Trusted Advisor console, schedule a job to export the cost optimization checks to a _csvfile_ Store the file in an S3 bucket every 2 weeks.fs20

Correct Answer: B
Section:
Explanation:
https://docs.aws.amazon.com/compute-optimizer/latest/APIReference/API_ExportLambdaFunctionRecommendations.html

QUESTION 48
The company needs to determine which costs on the monthly AWS bill are attributable to each application or team. The company also must be able to create reports to compare costs from the last 12 months and to help forecast costs for the next 12 months. A solutions architect must recommend an AWS Billing and Cost Management solution that provides these cost reports.
Which combination of actions will meet these requirements? (Select THREE.)

A. Activate the user-defined cost allocation tags that represent the application and the team.

VCEûp CEplus IT Certification Exams - Questions & Answers | VCEplus.io CEplus VCEûp

VCEplus

B. Activate the AWS generated cost allocation tags that represent the application and the team.
C. Create a cost category for each application in Billing and Cost Management.
D. Activate IAM access to Billing and Cost Management.
E. Create a cost budget.
F. Enable Cost Explorer.

Correct Answer: A, C, F
Section:
Explanation:
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/manage-cost-categories.html
https://aws.amazon.com/premiumsupport/knowledge-center/cost-explorer-analyze-spending-and-usage/
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/manage-cost-categories.html https://docs.aws.amazon.com/cost-management/latest/userguide/ce-enable.html
The best combination of actions to meet the company's requirements is Options A, C, and F.
Option A involves activating the user-defined cost allocation tags that represent the application and the team. This will allow the company to assign costs to different applications or teams, and will allow them to be tracked in the monthly AWS bill.
Option C involves creating a cost category for each application in Billing and Cost Management. This will allow the company to easily identify and compare costs across different applications and teams.
Option F involves enabling Cost Explorer. This will allow the company to view the costs of their AWS resources over the last 12 months and to create forecasts for the next 12 months.
These recommendations are in line with the official Amazon Textbook and Resources for the AWS Certified Solutions Architect - Professional certification. In particular, the book states that ''You can use cost allocation tags to group your costs by application, team, or other categories'' (Source: https://d1.awsstatic.com/training-and-certification/docs-sa-pro/AWS_Certified_Solutions_Architect_Professional_Exam_Guide_EN_v1.5.pdf). Additionally, the book states that ''Cost Explorer enables you to view the costs of your AWS resources over the last 12 months and to create forecasts for the next 12 months'' (Source: https://d1.awsstatic.com/training-and-certification/docs-sa-pro/AWS_Certified_Solutions_Architect_Professional_Exam_Guide_EN_v1.5.pdf).

QUESTION 49
An AWS customer has a web application that runs on premises. The web application fetches data from a third-party API that is behind a firewall. The third party accepts only one public CIDR block in each client's allow list.
The customer wants to migrate their web application to the AWS Cloud. The application will be hosted on a set of Amazon EC2 instances behind an Application Load Balancer (ALB) in a VPC. The ALB is located in public subnets. The EC2 instances are located in private subnets. NAT gateways provide internet access to the private subnets.
How should a solutions architect ensure that the web application can continue to call the third-parly API after the migration?

A. Associate a block of customer-owned public IP addresses to the VPC. Enable public IP addressing for public subnets in the VPC.
B. Register a block of customer-owned public IP addresses in the AWS account. Create Elastic IP addresses from the address block and assign them lo the NAT gateways in the VPC.
C. Create Elastic IP addresses from the block of customer-owned IP addresses. Assign the static Elastic IP addresses to the ALB.
D. Register a block of customer-owned public IP addresses in the AWS account. Set up AWS Global Accelerator to use Elastic IP addresses from the address block. Set the ALB as the accelerator endpoint.

Correct Answer: B
Section:
Explanation:
When EC2 instances reach third-party API through internet, their privates IP addresses will be masked by NAT Gateway public IP address.
https://aws.amazon.com/blogs/networking-and-content-delivery/introducing-bring-your-own-ip-byoip-for-amazon-vpc/

QUESTION 50

A company with several AWS accounts is using AWS Organizations and service control policies (SCPs). An Administrator created the following SCP and has attached it to an organizational unit (OU) that contains AWS account 1111-1111-1111:

VCEûp VCEplus IT Certification Exams - Questions & Answers | VCEplus.io VCEplus VCEûp

VCEplus

```
{
   "Version": "2012-10-17",
   "Statement": [
      {
         "Sid": "AllowsAllActions",
         "Effect": "Allow",
         "Action": "*",
         "Resource": "*"
      },
      {
         "Sid": "DenyCloudTrail",
         "Effect": "Deny",
         "Action": "cloudtrail:*",
         "Resource": "*"
      }
   ]
}
```

Developers working in account 1111-1111-1111 complain that they cannot create Amazon S3 buckets. How should the Administrator address this problem?

A. Add s3:CreateBucket with Allow effect to the SCP.
B. Remove the account from the OU, and attach the SCP directly to account 1111-1111-1111.
C. Instruct the Developers to add Amazon S3 permissions to their IAM entities.
D. Remove the SCP from account 1111-1111-1111.

Correct Answer: C
Section:
Explanation:
However A's explanation is incorrect - https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html
'SCPs are similar to AWS Identity and Access Management (IAM) permission policies and use almost the same syntax. However, an SCP never grants permissions.'
SCPs alone are not sufficient to granting permissions to the accounts in your organization. No permissions are granted by an SCP. An SCP defines a guardrail, or sets limits, on the actions that the account's administrator can delegate to the IAM users and roles in the affected accounts. The administrator must still attach identity-based or resource-based policies to IAM users or roles, or to the resources in your accounts to actually grant permissions. The effective permissions are the logical intersection between what is allowed by the SCP and what is allowed by the IAM and resource-based policies.

QUESTION 51
A company has a legacy monolithic application that is critical to the company's business. The company hosts the application on an Amazon EC2 instance that runs Amazon Linux 2. The company's application team receives a directive from the legal department to back up the data from the instance's encrypted Amazon
Elastic Block Store (Amazon EBS) volume to an Amazon S3 bucket. The application team does not have the administrative SSH key pair for the instance. The application must continue to serve the users.
Which solution will meet these requirements?

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

Correct Answer: C
Section:
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

VCEplus

the backup is performed on the EBS volume level.

QUESTION 52
A solutions architect needs to copy data from an Amazon S3 bucket m an AWS account to a new S3 bucket in a new AWS account. The solutions architect must implement a solution that uses the AWS CLI.
Which combination of steps will successfully copy the data? (Choose three.)

A. Create a bucket policy to allow the source bucket to list its contents and to put objects and set object ACLs in the destination bucket. Attach the bucket policy to the destination bucket.
B. Create a bucket policy to allow a user In the destination account to list the source bucket's contents and read the source bucket's objects. Attach the bucket policy to the source bucket.
C. Create an IAM policy in the source account. Configure the policy to allow a user In the source account to list contents and get objects In the source bucket, and to list contents, put objects, and set object ACLs in the destination bucket. Attach the policy to the user _
D. Create an IAM policy in the destination account. Configure the policy to allow a user In the destination account to list contents and get objects In the source bucket, and to list contents, put objects, and set objectACLs in the destination bucket. Attach the policy to the user.
E. Run the aws s3 sync command as a user in the source account. Specify' the source and destination buckets to copy the data.
F. Run the aws s3 sync command as a user in the destination account. Specify' the source and destination buckets to copy the data.

Correct Answer: B, D, F
Section:
Explanation:
Step B is necessary so that the user in the destination account has the necessary permissions to access the source bucket and list its contents, read its objects. Step D is needed so that the user in the destination account has the necessary permissions to access the destination bucket and list contents, put objects, and set object ACLs Step F is necessary because the aws s3 sync command needs to be run using the IAM user credentials from the destination account, so that the objects will have the appropriate permissions for the user in the destination account once they are copied.

QUESTION 53
A company built an application based on AWS Lambda deployed in an AWS CloudFormation stack. The last production release of the web application introduced an issue that resulted in an outage lasting several minutes. A solutions architect must adjust the deployment process to support a canary release.
Which solution will meet these requirements?

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

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

QUESTION 54
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 group. 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.
B. Migrate the SFTP server to AWS Transfer for SFTP. Update the DNS record sftp.example.com in Route 53 to point to the server endpoint hostname.

VCEplus

C. Migrate the SFTP server to a file gateway in AWS Storage Gateway. Update the DNS record sflp.example.com in Route 53 to point to the file gateway endpoint.
D. 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.

Correct Answer: B
Section:
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_ls

QUESTION 55
A company wants to migrate an application to Amazon EC2 from VMware Infrastructure that runs in an on-premises data center. A solutions architect must preserve the software and configuration settings during the migration.
What should the solutions architect do to meet these requirements?

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

www.VCEplus.io

Correct Answer: D
Section:
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/

QUESTION 56
A video processing company has an application that downloads images from an Amazon S3 bucket, processes the images, stores a transformed image in a second S3 bucket, and updates metadata about the image in an Amazon DynamoDB table. The application is written in Node.js and runs by using an AWS Lambda function. The Lambda function is invoked when a new image is uploaded to Amazon S3.
The application ran without incident for a while. However, the size of the images has grown significantly. The Lambda function is now failing frequently with timeout errors. The function timeout is set to its maximum value. A solutions architect needs to refactor the application's architecture to prevent invocation failures. The company does not want to manage the underlying infrastructure.
Which combination of steps should the solutions architect take to meet these requirements? (Choose two.)

A. Modify the application deployment by building a Docker image that contains the application code. Publish the image to Amazon Elastic Container Registry (Amazon ECR).
B. Create a new Amazon Elastic Container Service (Amazon ECS) task definition with a compatibility type of AWS Fargate. Configure the task definition to use the new image in Amazon Elastic Container Registry (Amazon ECR). Adjust the Lambda function to invoke an ECS task by using the ECS task definition when a new file arrives in Amazon S3.
C. Create an AWS Step Functions state machine with a Parallel state to invoke the Lambda function. Increase the provisioned concurrency of the Lambda function.
D. Create a new Amazon Elastic Container Service (Amazon ECS) task definition with a compatibility type of Amazon EC2. Configure the task definition to use the new image in Amazon Elastic Container Registry (Amazon ECR). Adjust the Lambda function to invoke an ECS task by using the ECS task definition when a new file arrives in Amazon S3.

CEplus

E. Modify the application to store images on Amazon Elastic File System (Amazon EFS) and to store metadata on an Amazon RDS DB instance. Adjust the Lambda function to mount the EFS file share.

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

QUESTION 57
A company has an organization in AWS Organizations. The company is using AWS Control Tower to deploy a landing zone for the organization. The company wants to implement governance and policy enforcement. The company must implement a policy that will detect Amazon RDS DB instances that are not encrypted at rest in the company's production OU.
Which solution will meet this requirement?

A. Turn on mandatory guardrails in AWS Control Tower. Apply the mandatory guardrails to the production OU.
B. Enable the appropriate guardrail from the list of strongly recommended guardrails in AWS Control Tower. Apply the guardrail to the production OU.
C. Use AWS Config to create a new mandatory guardrail. Apply the rule to all accounts in the production OU.
D. Create a custom SCP in AWS Control Tower. Apply the SCP to the production OU.

Correct Answer: B
Section:
Explanation:
AWS Control Tower provides a set of 'strongly recommended guardrails' that can be enabled to implement governance and policy enforcement. One of these guardrails is 'Encrypt Amazon RDS instances' which will detect RDS DB instances that are not encrypted at rest. By enabling this guardrail and applying it to the production OU, the company will be able to enforce encryption for RDS instances in the production environment.

QUESTION 58
A company recently completed the migration from an on-premises data center to the AWS Cloud by using a replatforming strategy. One of the migrated servers is running a legacy Simple Mail Transfer Protocol (SMTP) service that a critical application relies upon. The application sends outbound email messages to the company's customers. The legacy SMTP server does not support TLS encryption and uses TCP port 25. The application can use SMTP only.
The company decides to use Amazon Simple Email Service (Amazon SES) and to decommission the legacy SMTP server. The company has created and validated the SES domain. The company has lifted the SES limits.
What should the company do to modify the application to send email messages from Amazon SES?

A. Configure the application to connect to Amazon SES by using TLS Wrapper. Create an IAM role that has ses:SendEmail and ses:SendRawEmail permissions. Attach the IAM role to an Amazon EC2 instance.
B. Configure the application to connect to Amazon SES by using STARTTLS. Obtain Amazon SES SMTP credentials. Use the credentials to authenticate with Amazon SES.
C. Configure the application to use the SES API to send email messages. Create an IAM role that has ses:SendEmail and ses:SendRawEmail permissions. Use the IAM role as a service role for Amazon SES.
D. Configure the application to use AWS SDKs to send email messages. Create an IAM user for Amazon SES. Generate API access keys. Use the access keys to authenticate with Amazon SES.

Correct Answer: B
Section:
Explanation:
To set up a STARTTLS connection, the SMTP client connects to the Amazon SES SMTP endpoint on port 25, 587, or 2587, issues an EHLO command, and waits for the server to announce that it supports the STARTTLS SMTP extension. The client then issues the STARTTLS command, initiating TLS negotiation. When negotiation is complete, the client issues an EHLO command over the new encrypted connection, and the SMTP session proceeds normally To set up a TLS Wrapper connection, the SMTP client connects to the Amazon SES SMTP endpoint on port 465 or 2465. The server presents its certificate, the client issues an EHLO command, and the SMTP session proceeds normally.
https://docs.aws.amazon.com/ses/latest/dg/smtp-connect.html

VCEûp CEplus IT Certification Exams - Questions & Answers | VCEplus.io CEplus VCEûp

CEplus

QUESTION 59
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 organization. Define tags and cost categories in the report. Create a table in Amazon Athena. Create an Amazon QuickSight dataset based on the Athena table. Share the dataset with the finance team.
B. Create an AWS Cost and Usage Report for the organization. Define tags and cost categories in the report. Create a specialized template in AWS Cost Explorer that the finance department will use to build reports.
C. Create an Amazon QuickSight dataset that receives spending information from the AWS Price List Query API. Share the dataset with the finance team.
D. Use the AWS Price List Query API to collect account spending information. Create a specialized template in AWS Cost Explorer that the finance department will use to build reports.

Correct Answer: A
Section:
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.

QUESTION 60
A company runs an loT platform on AWS loT sensors in various locations send data to the company's Node js API servers on Amazon EC2 instances running behind an Application Load Balancer The data is stored in an Amazon RDS MySQL DB instance that uses a 4 TB General Purpose SSD volume
The number of sensors the company has deployed in the field has increased over time and is expected to grow significantly The API servers are consistently overloaded and RDS metrics show high write latency
Which of the following steps together will resolve the issues permanently and enable growth as new sensors are provisioned, while keeping this platform cost-efficient? {Select TWO.)

A. Resize the MySQL General Purpose SSD storage to 6 TB to improve the volume's IOPS
B. Re-architect the database tier to use Amazon Aurora instead of an RDS MySQL DB instance and add read replicas
C. Leverage Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data
D. Use AWS X-Ray to analyze and debug application issues and add more API servers to match the load
E. Re-architect the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance

Correct Answer: C, E
Section:
Explanation:
Option C is correct because leveraging Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data resolves the issues permanently and enable growth as new sensors are provisioned. Amazon Kinesis Data Streams is a serverless streaming data service that simplifies the capture, processing, and storage of data streams at any scale. Kinesis Data Streams can handle any amount of streaming data and process data from hundreds of thousands of sources with very low latency. AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. Lambda can be triggered by Kinesis Data Streams events and process the data records in real time. Lambda can also scale automatically based on the incoming data volume.By using Kinesis Data Streams and Lambda, the company can reduce the load on the API servers and improve the performance and scalability of the data ingestion and processing layer3
Option E is correct because re-architecting the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance resolves the issues permanently and enable growth as new sensors are provisioned. Amazon DynamoDB is a fully managed key-value and document database that delivers single-digit millisecond performance at any scale. DynamoDB supports auto scaling, which automatically adjusts read and write capacity based on actual traffic patterns. DynamoDB also supports on-demand capacity mode, which instantly accommodates up to double the previous peak traffic on a table. By using DynamoDB instead of RDS MySQL DB instance, the company can eliminate high write latency and improve scalability and performance of the database tier.

QUESTION 61
A company is building an electronic document management system in which users upload their documents. The application stack is entirely serverless and runs on AWS in the eu-central-1 Region. The system

VCEûp CEplus IT Certification Exams - Questions & Answers | VCEplus.io CEplus VCEûp

VCEplus

includes a web application that uses an Amazon CloudFront distribution for delivery with Amazon S3 as the origin. The web application communicates with Amazon API Gateway Regional endpoints. The API Gateway APIs call AWS Lambda functions that store metadata in an Amazon Aurora Serverless database and put the documents into an S3 bucket.
The company is growing steadily and has completed a proof of concept with its largest customer. The company must improve latency outside of Europe.
Which combination of actions will meet these requirements? (Select TWO.)

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

Correct Answer: A, C
Section:
Explanation:
https://aws.amazon.com/global-accelerator/faqs/

QUESTION 62
An adventure company has launched a new feature on its mobile app. Users can use the feature to upload their hiking and ratting photos and videos anytime. The photos and videos are stored in Amazon S3 Standard storage in an S3 bucket and are served through Amazon CloudFront.
The company needs to optimize the cost of the storage. A solutions architect discovers that most of the uploaded photos and videos are accessed infrequently after 30 days. However, some of the uploaded photos and videos are accessed frequently after 30 days. The solutions architect needs to implement a solution that maintains millisecond retrieval availability of the photos and videos at the lowest possible cost.
Which solution will meet these requirements?

A. Configure S3 Intelligent-Tiering on the S3 bucket.
B. Configure an S3 Lifecycle policy to transition image objects and video objects from S3 Standard to S3 Glacier Deep Archive after 30 days.
C. Replace Amazon S3 with an Amazon Elastic File System (Amazon EFS) file system that is mounted on Amazon EC2 instances.
D. Add a Cache-Control: max-age header to the S3 image objects and S3 video objects. Set the header to 30 days.

Correct Answer: B
Section:
Explanation:
Amazon S3 Intelligent-Tiering is a storage class that automatically moves objects between two access tiers based on changing access patterns. Objects that are accessed frequently are stored in the frequent access tier and objects that are accessed infrequently are stored in the infrequent access tier. This allows for cost optimization without requiring manual intervention. This makes it an ideal solution for the scenario described, as it can automatically move objects that are infrequently accessed after 30 days to a lower-cost storage tier while still maintaining millisecond retrieval availability.

QUESTION 63
A company uses Amazon S3 to store files and images in a variety of storage classes. The company's S3 costs have increased substantially during the past year.
A solutions architect needs to review data trends for the past 12 months and identity the appropriate storage class for the objects.
Which solution will meet these requirements?

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

Correct Answer: B

VCEûp VCEplus IT Certification Exams - Questions & Answers | VCEplus.io VCEplus VCEûp

CEplus

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

QUESTION 64
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

Correct Answer: C
Section:
Explanation:
https://aws.amazon.com/blogs/aws/new-use-aws-cloudformation-stacksets-for-multiple-accounts-in-an-aws-organization/
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.

QUESTION 65
A start up company hosts a fleet of Amazon EC2 instances in private subnets using the latest Amazon Linux 2 AMI. The company's engineers rely heavily on SSH access to the instances for troubleshooting.
The company's existing architecture includes the following:
* A VPC with private and public subnets, and a NAT gateway
* Site-to-Site VPN for connectivity with the on-premises environment
* EC2 security groups with direct SSH access from the on-premises environment
The company needs to increase security controls around SSH access and provide auditing of commands executed by the engineers.
Which strategy should a solutions architect use?

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

Correct Answer: D
Section:
Explanation:
Allows client machines to be able to connect to Session Manager using the AWS CLI instead of going through the AWS EC2 or AWS Server Manager console. https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager-working-with-install-plugin.html https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager-working-with-install-plugin.html#:~:text=aws%20ssm%20start%2Dsession%20%2D%2Dtarget%20instance%2Did

VCEûp CEplus IT Certification Exams - Questions & Answers | VCEplus.io CEplus VCEûp

VCEplus

QUESTION 66
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 limit. Apply the SCP to the developer accounts.
B. Use AWS Budgets to create a fixed monthly budget for each developer's account as part of the account creation process.
C. Create an SCP to deny access to costly services and components. Apply the SCP to the developer accounts.
D. Create an IAM policy to deny access to costly services and components. Apply the IAM policy to the developer accounts.
E. Create an AWS Budgets alert action to terminate services when the budgeted amount is reached. Configure the action to terminate all services.
F. Create an AWS Budgets alert action to send an Amazon Simple Notification Service (Amazon SNS) notification when the budgeted amount is reached. Invoke an AWS Lambda function to terminate all services.

Correct Answer: B, C, F
Section:
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.

QUESTION 67
A company is migrating some of its applications to AWS. The company wants to migrate and modernize the applications quickly after it finalizes networking and security strategies. The company has set up an AWS Direct Connection connection in a central network account.
The company expects to have hundreds of AWS accounts and VPCs in the near future. The corporate network must be able to access the resources on AWS seamlessly and also must be able to communicate with all the VPCs. The company also wants to route its cloud resources to the internet through its on-premises data center.
Which combination of steps will meet these requirements? (Choose three.)

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

CEplus

Correct Answer: B, D, F
Section:
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.

QUESTION 68
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 action. Attach the SCP to each OU of the organization.
E. Ensure that all AWS accounts are part of an organization in AWS Organizations that uses the consolidated billing feature.

Correct Answer: A, D
Section:
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

CEplus

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

QUESTION 69
A company is running a critical application that uses an Amazon RDS for MySQL database to store data. The RDS DB instance is deployed in Multi-AZ mode.
A recent RDS database failover test caused a 40-second outage to the application A solutions architect needs to design a solution to reduce the outage time to less than 20 seconds.
Which combination of steps should the solutions architect take to meet these requirements? (Select THREE.)

A. Use Amazon ElastiCache for Memcached in front of the database

B. Use Amazon ElastiCache for Redis in front of the database.

C. Use RDS Proxy in front of the database

D. Migrate the database to Amazon Aurora MySQL

E. Create an Amazon Aurora Replica

F. Create an RDS for MySQL read replica

Correct Answer: C, D, E
Section:
Explanation:
Migrate the database to Amazon Aurora MySQL. - Create an Amazon Aurora Replica. - Use RDS Proxy in front of the database. - These options are correct because they address the requirement of reducing the failover time to less than 20 seconds. Migrating to Amazon Aurora MySQL and creating an Aurora replica can reduce the failover time to less than 20 seconds. Aurora has a built-in, fault-tolerant storage system that can automatically detect and repair failures. Additionally, Aurora has a feature called 'Aurora Global Database' which allows you to create read-only replicas across multiple AWS regions which can further help to reduce the failover time. Creating an Aurora replica can also help to reduce the failover time as it can take over as the primary DB instance in case of a failure. Using RDS proxy can also help to reduce the failover time as it can route the queries to the healthy DB instance, it also helps to balance the load across multiple DB instances.

QUESTION 70
An AWS partner company is building a service in AWS Organizations using Its organization named org. This service requires the partner company to have access to AWS resources in a customer account, which is in a separate organization named org2 The company must establish least privilege security access using an API or command line tool to the customer account
What is the MOST secure way to allow org1 to access resources h org2?

A. The customer should provide the partner company with their AWS account access keys to log in and perform the required tasks
B. The customer should create an IAM user and assign the required permissions to the IAM user The customer should then provide the credentials to the partner company to log In and perform the required tasks.
C. The customer should create an IAM role and assign the required permissions to the IAM role. The partner company should then use the IAM rote's Amazon Resource Name (ARN) when requesting access to perform the required tasks
D. The customer should create an IAM rote and assign the required permissions to the IAM rote. The partner company should then use the IAM rote's Amazon Resource Name (ARN). Including the external ID in the IAM role's trust pokey, when requesting access to perform the required tasks

Correct Answer: C
Section:
Explanation:
https://docs.aws.amazon.com/IAM/latest/UserGuide/confused-deputy.html
This is the most secure way to allow org1 to access resources in org2 because it allows for least privilege security access. The customer should create an IAM role and assign the required permissions to the IAM role. The partner

CEplus

company should then use the IAM role's Amazon Resource Name (ARN) and include the external ID in the IAM role's trust policy when requesting access to perform the required tasks. This ensures that the partner company can only access the resources that it needs and only from the specific customer account.

QUESTION 71
A retail company has structured its AWS accounts to be part of an organization in AWS Organizations. The company has set up consolidated billing and has mapped its departments to the following OUs: Finance. Sales. Human Resources <HR). Marketing, and Operations. Each OU has multiple AWS accounts, one for each environment within a department. These environments are development, test, pre-production, and production.
The HR department is releasing a new system thai will launch in 3 months. In preparation, the HR department has purchased several Reserved Instances (RIs) in its production AWS account. The HR department will install the new application on this account. The HR department wants to make sure that other departments cannot share the RI discounts.
Which solution will meet these requirements?

A. In the AWS Billing and Cost Management console for the HR department's production account, turn off R1 sharing.
B. Remove the HR department's production AWS account from the organization. Add the account to the consolidating billing configuration only.
C. In the AWS Billing and Cost Management console, use the organization's management account to turn off R1 sharing for the HR department's production AWS account.
D. Create an SCP in the organization to restrict access to the RIs. Apply the SCP to the OUs of the other departments.

Correct Answer: C
Section:
Explanation:
You can use the management account of the organization in AWS Billing and Cost Management console to turn off RI sharing for the HR department's production AWS account. This will prevent other departments from sharing the RI discounts and ensure that only the HR department can use the RIs purchased in their production account.

QUESTION 72
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 process. Use Session Manager to log in to an instance that is marked as unhealthy
B. Enable EC2 instance termination protection Use Session Manager to log In to an instance that is marked as unhealthy.
C. Set the termination policy to Oldestinstance on the Auto Scaling group. Use Session Manager to log in to an instance that is marked as unhealthy
D. Suspend the Auto Scaling group's Terminate process. Use Session Manager to log in to an instance that is marked as unhealthy

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

QUESTION 73
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 organization. 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

VCEplus

update the security policy in the Firewall Manager administrative account

B. Deploy an organization-wide AWS Config rule that requires all resources in the selected OUs to associate the AWS WAF rules. 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.

C. 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.

D. 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

Correct Answer: A
Section:
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

QUESTION 74
A solutions architect is auditing the security setup of an AWS Lambda function for a company. The Lambda function retrieves the latest changes from an Amazon Aurora database. The Lambda function and the database run in the same VPC. Lambda environment variables are providing the database credentials to the Lambda function.
The Lambda function aggregates data and makes the data available in an Amazon S3 bucket that is configured for server-side encryption with AWS KMS managed encryption keys (SSE-KMS). The data must not travel across the internet. If any database credentials become compromised, the company needs a solution that minimizes the impact of the compromise.
What should the solutions architect recommend to meet these requirements?

www.VCEplus.io

A. Enable IAM database authentication on the Aurora DB cluster. Change the IAM role for the Lambda function to allow the function to access the database by using IAM database authentication. Deploy a gateway VPC endpoint for Amazon S3 in the VPC.

B. Enable IAM database authentication on the Aurora DB cluster. Change the IAM role for the Lambda function to allow the function to access the database by using IAM database authentication. Enforce HTTPS on the connection to Amazon S3 during data transfers.

C. Save the database credentials in AWS Systems Manager Parameter Store. Set up password rotation on the credentials in Parameter Store. Change the IAM role for the Lambda function to allow the function to access Parameter Store. Modify the Lambda function to retrieve the credentials from Parameter Store. Deploy a gateway VPC endpoint for Amazon S3 in the VPC.

D. Save the database credentials in AWS Secrets Manager. Set up password rotation on the credentials in Secrets Manager. Change the IAM role for the Lambda function to allow the function to access Secrets Manager. Modify the Lambda function to retrieve the credentials Om Secrets Manager. Enforce HTTPS on the connection to Amazon S3 during data transfers.

Correct Answer: A
Section:
Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/UsingWithRDS.IAMDBAuth.html

QUESTION 75
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?

VCEplus

A. Create a desired-instance-type managed rule in AWS Config. Configure the rule with the instance types that are allowed. Attach the rule to an event to run each time a new EC2 instance is launched.
B. In the EC2 console, create a launch template that specifies the instance types that are allowed. Assign the launch template to the developers' IAM accounts.
C. Create a new IAM policy. Specify the instance types that are allowed. Attach the policy to an IAM group that contains the IAM accounts for the developers
D. Use EC2 Image Builder to create an image pipeline for the developers and assist them in the creation of a golden image.

Correct Answer: C
Section:
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

QUESTION 76
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.

www.VCEplus.io

Correct Answer: A, B, E
Section:
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.html

QUESTION 77
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

Correct Answer: A, D

VCEplus

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

QUESTION 78
A team collects and routes behavioral data for an entire company The company runs a Multi-AZ VPC environment with public subnets, private subnets, and in internet gateway Each public subnet also contains a NAT gateway Most of the company's applications read from and write to Amazon Kinesis Data Streams. Most of the workloads am in private subnets.
A solutions architect must review the infrastructure The solutions architect needs to reduce costs and maintain the function of the applications The solutions architect uses Cost Explorer and notices that the cost in the EC2-Other category is consistently high A further review shows that NatGateway-Bytes charges are increasing the cost in the EC2-Other category.
What should the solutions architect do to meet these requirements?

A. Enable VPC Flow Logs. Use Amazon Athena to analyze the logs for traffic that can be removed. Ensure that security groups are Mocking traffic that is responsible for high costs.
B. Add an interface VPC endpoint for Kinesis Data Streams to the VPC. Ensure that applications have the correct IAM permissions to use the interface VPC endpoint.
C. Enable VPC Flow Logs and Amazon Detective Review Detective findings for traffic that is not related to Kinesis Data Streams Configure security groups to block that traffic
D. Add an interface VPC endpoint for Kinesis Data Streams to the VPC. Ensure that the VPC endpoint policy allows traffic from the applications.

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

www.VCEplus.io

QUESTION 79
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 gateway. Create a private VIF from the DX-B connection into the same Direct Connect gateway for high availability. Associate both the eu-west-1 and us-east-1 transit gateways with the Direct Connect gateway. Peer the transit gateways with each other to support cross-Region routing.
B. Create a transit VIF from the DX-A connection into a Direct Connect gateway. Associate the eu-west-1 transit gateway with this Direct Connect gateway. Create a transit VIF from the DX-B connection into a separate Direct Connect gateway. Associate the us-east-1 transit gateway with this separate Direct Connect gateway. Peer the Direct Connect gateways with each other to support high availability and cross-Region routing.
C. Create a transit VIF from the DX-A connection into a Direct Connect gateway. Create a transit VIF from the DX-B connection into the same Direct Connect gateway for high availability. Associate both the eu-west-1 and us-east-1 transit gateways with this Direct Connect gateway. Configure the Direct Connect gateway to route traffic between the transit gateways.
D. Create a transit VIF from the DX-A connection into a Direct Connect gateway. Create a transit VIF from the DX-B connection into the same Direct Connect gateway for high availability. Associate both the eu-west-1 and us-east-1 transit gateways with this Direct Connect gateway. Peer the transit gateways with each other to support cross-Region routing.

Correct Answer: D
Section:
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

VCEûp VCEplus IT Certification Exams - Questions & Answers | VCEplus.io VCEplus VCEûp

VCEplus

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.

QUESTION 80
A company is running an application in the AWS Cloud. The company's security team must approve the creation of all new IAM users. When a new IAM user is created, all access for the user must be removed automatically. The security team must then receive a notification to approve the user. The company has a multi-Region AWS CloudTrail trail In the AWS account.
Which combination of steps will meet these requirements? (Select THREE.)

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

Correct Answer: A, D, E
Section:
Explanation:
https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/send-a-notification-when-an-iam-user-is-created.html

QUESTION 81
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 Tower. Enroll accounts and invite existing accounts into the resulting organization in AWS Organizations.
B. Enable AWS Security Hub in all accounts to manage cross-account access. Collect findings through AWS CloudTrail to force MFA login.
C. Create transit gateways and transit gateway VPC attachments in each account. Configure appropriate route tables.
D. Set up and enable AWS IAM Identity Center (AWS Single Sign-On). Create appropriate permission sets with required MFA for existing accounts.
E. Enable AWS Control Tower in all Recounts to manage routing between accounts. Collect findings through AWS CloudTrail to force MFA login.
F. Create IAM users and groups. Configure MFA for all users. Set up Amazon Cognito user pools and identity pools to manage access to accounts and between accounts.

Correct Answer: A, C, D
Section:
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.

QUESTION 82
A delivery company needs to migrate its third-party route planning application to AWS. The third party supplies a supported Docker image from a public registry. The image can run in as many containers as required to generate the route map.
The company has divided the delivery area into sections with supply hubs so that delivery drivers travel the shortest distance possible from the hubs to the customers. To reduce the time necessary to generate

CEplus

route maps, each section uses its own set of Docker containers with a custom configuration that processes orders only in the section's area.
The company needs the ability to allocate resources cost-effectively based on the number of running containers.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create an Amazon Elastic Kubernetes Service (Amazon EKS) cluster on Amazon EC2. Use the Amazon EKS CLI to launch the planning application in pods by using the -tags option to assign a custom tag to the pod.
B. Create an Amazon Elastic Kubernetes Service (Amazon EKS) cluster on AWS Fargate. Use the Amazon EKS CLI to launch the planning application. Use the AWS CLI tag-resource API call to assign a custom tag to the pod.
C. Create an Amazon Elastic Container Service (Amazon ECS) cluster on Amazon EC2. Use the AWS CLI with run-tasks set to true to launch the planning application by using the -tags option to assign a custom tag to the task.
D. Create an Amazon Elastic Container Service (Amazon ECS) cluster on AWS Fargate. Use the AWS CLI run-task command and set enableECSManagedTags to true to launch the planning application. Use the --tags option to assign a custom tag to the task.

Correct Answer: D
Section:
Explanation:
Amazon Elastic Container Service (ECS) on AWS Fargate is a fully managed service that allows you to run containers without having to manage the underlying infrastructure. When you launch tasks on Fargate, resources are automatically allocated based on the number of tasks running, which reduces the operational overhead.
Using ECS on Fargate allows you to assign custom tags to tasks using the --tags option in the run-task command, as described in the documentation: https://docs.aws.amazon.com/cli/latest/reference/ecs/run-task.html You can also set enableECSManagedTags to true, which allows the service to automatically add the cluster name and service name as tags. https://docs.aws.amazon.com/AmazonECS/latest/developerguide/task-placement-constraints.html#tag-based-scheduling

QUESTION 83
A software company hosts an application on AWS with resources in multiple AWS accounts and Regions. The application runs on a group of Amazon EC2 instances in an application VPC located in the us-east-1 Region with an IPv4 CIDR block of 10.10.0.0/16. In a different AWS account, a shared services VPC is located in the us-east-2 Region with an IPv4 CIDR block of 10.10.10.0/24. When a cloud engineer uses AWS CloudFormation to attempt to peer the application
VPC with the shared services VPC, an error message indicates a peering failure.
Which factors could cause this error? (Choose two.)

A. The IPv4 CIDR ranges of the two VPCs overlap
B. The VPCs are not in the same Region
C. One or both accounts do not have access to an Internet gateway
D. One of the VPCs was not shared through AWS Resource Access Manager
E. The IAM role in the peer accepter account does not have the correct permissions

Correct Answer: A, E
Section:
Explanation:
https://aws.amazon.com/about-aws/whats-new/2017/11/announcing-support-for-inter-region-vpc-peering/

QUESTION 84
A company's solutions architect is reviewing a new internally developed application in a sandbox AWS account The application uses an AWS Auto Scaling group of Amazon EC2 instances that have an IAM instance profile attached Part of the application logic creates and accesses secrets from AWS Secrets Manager The company has an AWS Lambda function that calls the application API to test the functionality The company also has created an AWS CloudTrail trail in the account
The application's developer has attached the SecretsManagerReadWnte AWS managed IAM policy to an IAM role The IAM role is associated with the instance profile that is attached to the EC2 instances The solutions architect has invoked the Lambda function for testing
The solutions architect must replace the SecretsManagerReadWnte policy with a new policy that provides least privilege access to the Secrets Manager actions that the application requires

VCEûp CEplus IT Certification Exams - Questions & Answers | VCEplus.io CEplus VCEûp

CEplus

What is the MOST operationally efficient solution that meets these requirements?

A. Generate a policy based on CloudTrail events for the IAM role Use the generated policy output to create a new IAM policy Use the newly generated IAM policy to replace the SecretsManagerReadWnte policy that is attached to the IAM role
B. Create an analyzer in AWS Identity and Access Management Access Analyzer Use the IAM role's Access Advisor findings to create a new IAM policy Use the newly created IAM policy to replace the SecretsManagerReadWnte policy that is attached to the IAM role
C. Use the aws cloudtrail lookup-events AWS CLI command to filter and export CloudTrail events that are related to Secrets Manager Use a new IAM policy that contains the actions from CloudTrail to replace the SecretsManagerReadWnte policy that is attached to the IAM role
D. Use the IAM policy simulator to generate an IAM policy for the IAM role Use the newly generated IAM policy to replace the SecretsManagerReadWnte policy that is attached to the IAM role

Correct Answer: B
Section:
Explanation:
The IAM policy simulator will generate a policy that contains only the necessary permissions for the application to access Secrets Manager, providing the least privilege necessary to get the job done. This is the most efficient solution as it will not require additional steps such as analyzing CloudTrail events or manually creating and testing an IAM policy.
You can use the IAM policy simulator to generate an IAM policy for an IAM role by specifying the role and the API actions and resources that the application or service requires. The simulator will then generate an IAM policy that grants the least privilege access to those actions and resources.
Once you have generated an IAM policy using the simulator, you can replace the existing SecretsManagerReadWnte policy that is attached to the IAM role with the newly generated policy. This will ensure that the application or service has the least privilege access to the Secrets Manager actions that it requires.
You can access the IAM policy simulator through the IAM console, AWS CLI, and AWS SDKs. Here is the link for more information:
https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_simulator.html

QUESTION 85
A solutions architect must analyze a company's Amazon EC2 Instances and Amazon Elastic Block Store (Amazon EBS) volumes to determine whether the company is using resources efficiently The company is running several large, high-memory EC2 instances lo host database dusters that are deployed in active/passive configurations The utilization of these EC2 instances varies by the applications that use the databases, and the company has not identified a pattern
The solutions architect must analyze the environment and take action based on the findings.
Which solution meets these requirements MOST cost-effectively?

A. Create a dashboard by using AWS Systems Manager OpsConter Configure visualizations tor Amazon CloudWatch metrics that are associated with the EC2 instances and their EBS volumes Review the dashboard periodically and identify usage patterns Right size the EC2 instances based on the peaks in the metrics
B. Turn on Amazon CloudWatch detailed monitoring for the EC2 instances and their EBS volumes Create and review a dashboard that is based on the metrics Identify usage patterns Right size the FC? instances based on the peaks In the metrics
C. Install the Amazon CloudWatch agent on each of the EC2 Instances Turn on AWS Compute Optimizer, and let it run for at least 12 hours Review the recommendations from Compute Optimizer, and right size the EC2 instances as directed
D. Sign up for the AWS Enterprise Support plan Turn on AWS Trusted Advisor Wait 12 hours Review the recommendations from Trusted Advisor, and rightsize the EC2 instances as directed

Correct Answer: C
Section:
Explanation:
(https://aws.amazon.com/compute-optimizer/pricing/ , https://aws.amazon.com/systems-manager/pricing/).
https://aws.amazon.com/compute-optimizer/

QUESTION 86
A company uses AWS Organizations for a multi-account setup in the AWS Cloud. The company uses AWS Control Tower for governance and uses AWS Transit Gateway for VPC connectivity across accounts.
In an AWS application account, the company's application team has deployed a web application that uses AWS Lambda and Amazon RDS. The company's database administrators have a separate DBA account and use the account to centrally manage all the databases across the organization. The database administrators use an Amazon EC2 instance that is deployed in the DBA account to access an RDS database that is

VCEûp CEplus IT Certification Exams - Questions & Answers | VCEplus.io CEplus VCEûp

CEplus

deployed in the application account.
The application team has stored the database credentials as secrets in AWS Secrets Manager in the application account. The application team is manually sharing the secrets with the database administrators. The secrets are encrypted by the default AWS managed key for Secrets Manager in the application account. A solutions architect needs to implement a solution that gives the database administrators access to the database and eliminates the need to manually share the secrets.
Which solution will meet these requirements?

A. Use AWS Resource Access Manager (AWS RAM) to share the secrets from the application account with the DBA account. In the DBA account, create an IAM role that is named DBA-Admin. Grant the role the required permissions to access the shared secrets. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.
B. In the application account, create an IAM role that is named DBA-Secret. Grant the role the required permissions to access the secrets. In the DBA account, create an IAM role that is named DBA-Admin. Grant the DBA-Admin role the required permissions to assume the DBA-Secret role in the application account. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.
C. In the DBA account, create an IAM role that is named DBA-Admin. Grant the role the required permissions to access the secrets and the default AWS managed key in the application account. In the application account, attach resource-based policies to the key to allow access from the DBA account. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.
D. In the DBA account, create an IAM role that is named DBA-Admin. Grant the role the required permissions to access the secrets in the application account. Attach an SCP to the application account to allow access to the secrets from the DBA account. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.

Correct Answer: B
Section:
Explanation:
Option B is correct because creating an IAM role in the application account that has permissions to access the secrets and creating an IAM role in the DBA account that has permissions to assume the role in the application account eliminates the need to manually share the secrets. This approach uses cross-account IAM roles to grant access to the secrets in the application account.The database administrators can assume the role in the application account from their EC2 instance in the DBA account and retrieve the secrets without having to store them locally or share them manually2

QUESTION 87
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 ec2Reoon condition key to restrict access to all AWS Regions except ap-northeast-1. Apply the SCP to the root OU. the DataOps OU. and the Research OU.
E. Create an SCP Use the ec2:InstanceType condition key to restrict access to specific instance types Apply the SCP to the DataOps OU.

Correct Answer: C, E
Section:
Explanation:
https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_examples_aws_deny-requested-region.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_ec2.html

QUESTION 88
A company runs a serverless application in a single AWS Region. The application accesses external URLs and extracts metadata from those sites. The company uses an Amazon Simple Notification Service (Amazon SNS) topic to publish URLs to an Amazon Simple Queue Service (Amazon SQS) queue An AWS Lambda function uses the queue as an event source and processes the URLs from the queue Results are saved to an Amazon S3 bucket
The company wants to process each URL other Regions to compare possible differences in site localization URLs must be published from the existing Region. Results must be written to the existing S3 bucket in the current Region.

VCEûp CEplus IT Certification Exams - Questions & Answers | VCEplus.io CEplus VCEûp

VCEplus

Which combination of changes will produce multi-Region deployment that meets these requirements? (Select TWO.)

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

Correct Answer: A, C
Section:
Explanation:
https://docs.aws.amazon.com/sns/latest/dg/sns-cross-region-delivery.html

QUESTION 89
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 application. Use Amazon CloudWatch Logs to invoke the Lambda function every 4 hours.
B. Use AWS Batch to run the application. Use an AWS Step Functions state machine to invoke the AWS Batch job every 4 hours.
C. Use AWS Fargate to run the application. Use Amazon EventBridge (Amazon CloudWatch Events) to invoke the Fargate task every 4 hours.
D. Use Amazon EC2 Spot Instances to run the application. Use AWS CodeDeploy to deploy and run the application every 4 hours.

www.VCEplus.io

Correct Answer: C
Section:
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-and-aws-batch/ makes sense only for HPC, this is a single instance that needs to be run

QUESTION 90
A company is creating a sequel for a popular online game. A large number of users from all over the world will play the game within the first week after launch. Currently, the game consists of the following components deployed in a single AWS Region:
* Amazon S3 bucket that stores game assets
* Amazon DynamoDB table that stores player scores
A solutions architect needs to design a multi-Region solution that will reduce latency improve reliability, and require the least effort to implement
What should the solutions architect do to meet these requirements?

A. Create an Amazon CloudFront distribution to serve assets from the S3 bucket Configure S3 Cross-Region Replication Create a new DynamoDB able in a new Region Use the new table as a replica target tor DynamoDB global tables.
B. Create an Amazon CloudFront distribution to serve assets from the S3 bucket. Configure S3 Same-Region Replication. Create a new DynamoDB able m a new Region. Configure asynchronous replication between the DynamoDB tables by using AWS Database Migration Service (AWS DMS) with change data capture (CDC)
C. Create another S3 bucket in a new Region and configure S3 Cross-Region Replication between the buckets Create an Amazon CloudFront distribution and configure origin failover with two origins accessing the S3 buckets in each Region. Configure DynamoDB global tables by enabling Amazon DynamoDB Streams, and add a replica table in a new Region.
D. Create another S3 bucket in the same Region, and configure S3 Same-Region Replication between the buckets- Create an Amazon CloudFront distribution and configure origin failover with two origin accessing the S3 buckets Create a new DynamoDB table m a new Region Use the new table as a replica target for DynamoDB global tables.

VCEplus

Correct Answer: C
Section:
Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/dynamodb-global-table-stream-lambda/?nc1=h_ls

QUESTION 91
A company has an on-premises website application that provides real estate information for potential renters and buyers. The website uses a Java backend and a NOSQL MongoDB database to store subscriber data.
The company needs to migrate the entire application to AWS with a similar structure. The application must be deployed for high availability, and the company cannot make changes to the application
Which solution will meet these requirements?

A. use an Amazon Aurora DB cluster as the database for the subscriber data. Deploy Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones for the Java backend application.

B. Use MongoDB on Amazon EC2 instances as the database for the subscriber data. Deploy EC2 instances in an Auto Scaling group in a single Availability Zone for the Java backend application.

C. Configure Amazon DocumentD3 (with MongoDB compatibility) with appropriately sized instances in multiple Availability Zones as the database for the subscriber data. Deploy Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones for the Java backend application.

D. Configure Amazon DocumentDB (with MongoDB compatibility) in on-demand capacity mode in multiple Availability Zones as the database for the subscriber data. Deploy Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones for the Java backend application.

Correct Answer: C
Section:
Explanation:
On-demand capacity mode is the function of Dynamodb. https://aws.amazon.com/blogs/news/running-spiky-workloads-and-optimizing-costs-by-more-than-90-using-amazon-dynamodb-on-demand-capacity-mode/
Amazon DocumentDB Elastic Clusters https://aws.amazon.com/blogs/news/announcing-amazon-documentdb-elastic-clusters/
Deploy Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones for the Java backend application. This will provide high availability and scalability, while allowing the company to retain the same database structure as the original application.

QUESTION 92
A digital marketing company has multiple AWS accounts that belong to various teams. The creative team uses an Amazon S3 bucket in its AWS account to securely store images and media files that are used as content for the company's marketing campaigns. The creative team wants to share the S3 bucket with the strategy team so that the strategy team can view the objects.
A solutions architect has created an IAM role that is named strategy_reviewer in the Strategy account. The solutions architect also has set up a custom AWS Key Management Service (AWS KMS) key in the Creative account and has associated the key with the S3 bucket. However, when users from the Strategy account assume the IAM role and try to access objects in the S3 bucket, they receive an Account.
The solutions architect must ensure that users in the Strategy account can access the S3 bucket. The solution must provide these users with only the minimum permissions that they need.
Which combination of steps should the solutions architect take to meet these requirements? (Select THREE.)

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

Correct Answer: A, C, F
Section:

VCEûp VCEplus IT Certification Exams - Questions & Answers | VCEplus.io VCEplus VCEûp

CEplus

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/cross-account-access-denied-error-s3/

QUESTION 93
A life sciences company is using a combination of open source tools to manage data analysis workflows and Docker containers running on servers in its on-premises data center to process genomics data Sequencing data is generated and stored on a local storage area network (SAN), and then the data is processed. The research and development teams are running into capacity issues and have decided to re-architect their genomics analysis platform on AWS to scale based on workload demands and reduce the turnaround time from weeks to days
The company has a high-speed AWS Direct Connect connection Sequencers will generate around 200 GB of data for each genome, and individual jobs can take several hours to process the data with ideal compute capacity. The end result will be stored in Amazon S3. The company is expecting 10-15 job requests each day
Which solution meets these requirements?

A. Use regularly scheduled AWS Snowball Edge devices to transfer the sequencing data into AWS When AWS receives the Snowball Edge device and the data is loaded into Amazon S3 use S3 events to trigger an AWS Lambda function to process the data
B. Use AWS Data Pipeline to transfer the sequencing data to Amazon S3 Use S3 events to trigger an Amazon EC2 Auto Scaling group to launch custom-AMI EC2 instances running the Docker containers to process the data
C. Use AWS DataSync to transfer the sequencing data to Amazon S3 Use S3 events to trigger an AWS Lambda function that starts an AWS Step Functions workflow Store the Docker images in Amazon Elastic Container Registry (Amazon ECR) and trigger AWS Batch to run the container and process the sequencing data
D. Use an AWS Storage Gateway file gateway to transfer the sequencing data to Amazon S3 Use S3 events to trigger an AWS Batch job that runs on Amazon EC2 instances running the Docker containers to process the data

Correct Answer: C
Section:
Explanation:
AWS DataSync can be used to transfer the sequencing data to Amazon S3, which is a more efficient and faster method than using Snowball Edge devices. Once the data is in S3, S3 events can trigger an AWS Lambda function that starts an AWS Step Functions workflow. The Docker images can be stored in Amazon Elastic Container Registry (Amazon ECR) and AWS Batch can be used to run the container and process the sequencing data.

QUESTION 94
A company runs a content management application on a single Windows Amazon EC2 instance in a development environment. The application reads and writes static content to a 2 TB Amazon Elastic Block Store (Amazon EBS) volume that is attached to the instance as the root device. The company plans to deploy this application in production as a highly available and fault-tolerant solution that runs on at least three EC2 instances across multiple Availability Zones.
A solutions architect must design a solution that joins all the instances that run the application to an Active Directory domain. The solution also must implement Windows ACLs to control access to file contents. The application always must maintain exactly the same content on all running instances at any given point in time.
Which solution will meet these requirements with the LEAST management overhead?

A. Create an Amazon Elastic File System (Amazon EFS) file share. Create an Auto Scaling group that extends across three Availability Zones and maintains a minimum size of three instances. Implement a user data script to install the application, join the instance to the AD domain, and mount the EFS file share.
B. Create a new AMI from the current EC2 instance that is running. Create an Amazon FSx for Lustre file system. Create an Auto Scaling group that extends across three Availability Zones and maintains a minimum size of three instances. Implement a user data script to join the instance to the AD domain and mount the FSx for Lustre file system.
C. Create an Amazon FSx for Windows File Server file system. Create an Auto Scaling group that extends across three Availability Zones and maintains a minimum size of three instances. Implement a user data script to install the application and mount the FSx for Windows File Server file system. Perform a seamless domain join to join the instance to the AD domain.
D. Create a new AMI from the current EC2 instance that is running. Create an Amazon Elastic File System (Amazon EFS) file system. Create an Auto Scaling group that extends across three Availability Zones and maintains a minimum size of three instances. Perform a seamless domain join to join the instance to the AD domain.

Correct Answer: C
Section:
Explanation:

VCEûp CEplus IT Certification Exams - Questions & Answers | VCEplus.io CEplus VCEûp

VCEplus

https://docs.aws.amazon.com/fsx/latest/WindowsGuide/what-is.html https://docs.aws.amazon.com/directoryservice/latest/admin-guide/ms_ad_join_instance.html

QUESTION 95
A software as a service (SaaS) based company provides a case management solution to customers A3 part of the solution. The company uses a standalone Simple Mail Transfer Protocol (SMTP) server to send email messages from an application. The application also stores an email template for acknowledgement email messages that populate customer data before the application sends the email message to the customer.
The company plans to migrate this messaging functionality to the AWS Cloud and needs to minimize operational overhead.
Which solution will meet these requirements MOST cost-effectively?

A. Set up an SMTP server on Amazon EC2 instances by using an AMI from the AWS Marketplace. Store the email template in an Amazon S3 bucket. Create an AWS Lambda function to retrieve the template from the S3 bucket and to merge the customer data from the application with the template. Use an SDK in the Lambda function to send the email message.
B. Set up Amazon Simple Email Service (Amazon SES) to send email messages. Store the email template in an Amazon S3 bucket. Create an AWS Lambda function to retrieve the template from the S3 bucket and to merge the customer data from the application with the template. Use an SDK in the Lambda function to send the email message.
C. Set up an SMTP server on Amazon EC2 instances by using an AMI from the AWS Marketplace. Store the email template in Amazon Simple Email Service (Amazon SES) with parameters for the customer data. Create an AWS Lambda function to call the SES template and to pass customer data to replace the parameters. Use the AWS Marketplace SMTP server to send the email message.
D. Set up Amazon Simple Email Service (Amazon SES) to send email messages. Store the email template on Amazon SES with parameters for the customer data. Create an AWS Lambda function to call the SendTemplatedEmail API operation and to pass customer data to replace the parameters and the email destination.

Correct Answer: D
Section:
Explanation:
In this solution, the company can use Amazon SES to send email messages, which will minimize operational overhead as SES is a fully managed service that handles sending and receiving email messages. The company can store the email template on Amazon SES with parameters for the customer data and use an AWS Lambda function to call the SendTemplatedEmail API operation, passing in the customer data to replace the parameters and the email destination. This solution eliminates the need to set up and manage an SMTP server on EC2 instances, which can be costly and time-consuming.

www.VCEplus.io

QUESTION 96
A company is processing videos in the AWS Cloud by using Amazon EC2 instances in an Auto Scaling group. It takes 30 minutes to process a video. Several EC2 instances scale in and out depending on the number of videos in an Amazon Simple Queue Service (Amazon SQS) queue.
The company has configured the SQS queue with a redrive policy that specifies a target dead-letter queue and a maxReceiveCount of 1. The company has set the visibility timeout for the SQS queue to 1 hour. The company has set up an Amazon CloudWatch alarm to notify the development team when there are messages in the dead-letter queue.
Several times during the day, the development team receives notification that messages are in the dead-letter queue and that videos have not been processed properly. An investigation finds no errors in the application logs.
How can the company solve this problem?

A. Turn on termination protection for the EC2 instances.
B. Update the visibility timeout for the SOS queue to 3 hours.
C. Configure scale-in protection for the instances during processing.
D. Update the redrive policy and set maxReceiveCount to 0.

Correct Answer: B
Section:
Explanation:
The best solution for this problem is to update the visibility timeout for the SQS queue to 3 hours. This is because when the visibility timeout is set to 1 hour, it means that if the EC2 instance doesn't process the message within an hour, it will be moved to the dead-letter queue. By increasing the visibility timeout to 3 hours, this should give the EC2 instance enough time to process the message before it gets moved to the dead-letter queue. Additionally, configuring scale-in protection for the EC2 instances during processing will help to ensure that the instances are not terminated while the messages are being processed.

QUESTION 97
A company has developed APIs that use Amazon API Gateway with Regional endpoints. The APIs call AWS Lambda functions that use API Gateway authentication mechanisms. After a design review, a solutions

VCEûp VCEplus IT Certification Exams - Questions & Answers | VCEplus.io VCEplus VCEûp

VCEplus

architect identifies a set of APIs that do not require public access.
The solutions architect must design a solution to make the set of APIs accessible only from a VPC. All APIs need to be called with an authenticated user.
Which solution will meet these requirements with the LEAST amount of effort?

A. Create an internal Application Load Balancer (ALB). Create a target group. Select the Lambda function to call. Use the ALB DNS name to call the API from the VPC.
B. Remove the DNS entry that is associated with the API in API Gateway. Create a hosted zone in Amazon Route 53. Create a CNAME record in the hosted zone. Update the API in API Gateway with the CNAME record. Use the CNAME record to call the API from the VPC.
C. Update the API endpoint from Regional to private in API Gateway. Create an interface VPC endpoint in the VPC. Create a resource policy, and attach it to the API. Use the VPC endpoint to call the API from the VPC.
D. Deploy the Lambda functions inside the VPC. Provision an EC2 instance, and install an Apache server. From the Apache server, call the Lambda functions. Use the internal CNAME record of the EC2 instance to call the API from the VPC.

Correct Answer: C
Section:
Explanation:
This solution requires the least amount of effort as it only requires to update the API endpoint to private in API Gateway and create an interface VPC endpoint. Then create a resource policy and attach it to the API. This will make the API only accessible from the VPC and still keep the authentication mechanism intact.
Reference:
https://aws.amazon.com/premiumsupport/knowledge-center/private-api-gateway-vpc-endpoint/
https://aws.amazon.com/api-gateway/features/

QUESTION 98
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 distribution. Use Lambda@Edge to modify requests from North America to use the new origin.

Correct Answer: B, D
Section:
Explanation:
https://aws.amazon.com/about-aws/whats-new/2016/04/transfer-files-into-amazon-s3-up-to-300-percent-faster/

QUESTION 99
A solutions architect is investigating an issue in which a company cannot establish new sessions in Amazon Workspaces. An initial analysis indicates that the issue involves user profiles. The Amazon Workspaces environment is configured to use Amazon FSx for Windows File Server as the profile share storage. The FSx for Windows File Server file system is configured with 10 TB of storage.
The solutions architect discovers that the file system has reached its maximum capacity. The solutions architect must ensure that users can regain access. The solution also must prevent the problem from occurring again.
Which solution will meet these requirements?

A. Remove old user profiles to create space. Migrate the user profiles to an Amazon FSx for Lustre file system.

VCEûp VCEplus IT Certification Exams - Questions & Answers | VCEplus.io VCEplus VCEûp

VCEplus

B. Increase capacity by using the update-file-system command. Implement an Amazon CloudWatch metric that monitors free space. Use Amazon EventBridge to invoke an AWS Lambda function to increase capacity as required.
C. Monitor the file system by using the FreeStorageCapacity metric in Amazon CloudWatch. Use AWS Step Functions to increase the capacity as required.
D. Remove old user profiles to create space. Create an additional FSx for Windows File Server file system. Update the user profile redirection for 50% of the users to use the new file system.

Correct Answer: B
Section:
Explanation:
It can prevent the issue from happening again by monitoring the file system with the FreeStorageCapacity metric in Amazon CloudWatch and using Amazon EventBridge to invoke an AWS Lambda function to increase the capacity as required. This ensures that the file system always has enough free space to store user profiles and avoids reaching maximum capacity.

QUESTION 100
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 instance. Create an Auto Scaling group of EC2 instances behind an Application Load Balancer. Configure the Auto Scaling group to have a minimum of three instances.
B. 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 instance. Point the EC2 instance to the new path for file processing.
C. 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 function. Configure the Lambda function to add the metadata and update the delivery system.
D. 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 function. Configure the Lambda function to add the metadata and update the delivery system.

Correct Answer: C
Section:
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.

QUESTION 101
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.

VCEûp VCEplus IT Certification Exams - Questions & Answers | VCEplus.io VCEplus VCEûp

CEplus

Correct Answer: A
Section:
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.

QUESTION 102
A financial services company receives a regular data feed from its credit card servicing partner Approximately 5.000 records are sent every 15 minutes in plaintext, delivered over HTTPS directly into an Amazon S3 bucket with server-side encryption. This feed contains sensitive credit card primary account number (PAN) data The company needs to automatically mask the PAN before sending the data to another S3 bucket for additional internal processing. The company also needs to remove and merge specific fields, and then transform the record into JSON format Additionally, extra feeds are likely to be added in the future, so any design needs to be easily expandable.
Which solutions will meet these requirements?

A. Trigger an AWS Lambda function on file delivery that extracts each record and writes it to an Amazon SQS queue. Trigger another Lambda function when new messages arrive in the SQS queue to process the records, writing the results to a temporary location in Amazon S3. Trigger a final Lambda function once the SQS queue is empty to transform the records into JSON format and send the results to another S3 bucket for internal processing.
B. Trigger an AWS Lambda function on file delivery that extracts each record and writes it to an Amazon SQS queue. Configure an AWS Fargate container application to automatically scale to a single instance when the SQS queue contains messages. Have the application process each record, and transform the record into JSON format. When the queue is empty, send the results to another S3 bucket for internal processing and scale down the AWS Fargate instance.
C. Create an AWS Glue crawler and custom classifier based on the data feed formats and build a table definition to match. Trigger an AWS Lambda function on file delivery to start an AWS Glue ETL job to transform the entire record according to the processing and transformation requirements. Define the output format as JSON. Once complete, have the ETL job send the results to another S3 bucket for internal processing.
D. Create an AWS Glue crawler and custom classifier based upon the data feed formats and build a table definition to match. Perform an Amazon Athena query on file delivery to start an Amazon EMR ETL job to transform the entire record according to the processing and transformation requirements. Define the output format as JSON. Once complete, send the results to another S3 bucket for internal processing and scale down the EMR cluster.

www.VCEplus.io

Correct Answer: C
Section:
Explanation:
You can use a Glue crawler to populate the AWS Glue Data Catalog with tables. The Lambda function can be triggered using S3 event notifications when object create events occur. The Lambda function will then trigger the Glue ETL job to transform the records masking the sensitive data and modifying the output format to JSON. This solution meets all requirements.

QUESTION 103
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 servers. Set up replication for all servers. Launch test instances for regular drills. Cut over to the test instances to fail over the workload in the case of a failure event.
B. Install the AWS Replication Agent on the source servers, including the MySQL servers. Initialize AWS Elastic Disaster Recovery in the target AWS Region. Define the launch settings. Frequently perform failover and fallback from the most recent point in time.
C. Create AWS Database Migration Service (AWS DMS) replication servers and a target Amazon Aurora MySQL DB cluster to host the database. Create a DMS replication task to copy the existing data to the target DB cluster. Create a local AWS Schema Conversion Tool (AWS SCT) change data capture (CDC) task to keep the data synchronized. Install the rest of the software on EC2 instances by starting with a compatible base AMI.
D. Deploy an AWS Storage Gateway Volume Gateway on premises. Mount volumes on all on-premises servers. Install the application and the MySQL database on the new volumes. Take regular snapshots. Install all the software on EC2 Instances by starting with a compatible base AMI. Launch a Volume Gateway on an EC2 instance. Restore the volumes from the latest snapshot. Mount the new volumes on the EC2 instances in the case of a failure event.

VCEplus

Correct Answer: B
Section:
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

QUESTION 104
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 account. Assign a unique external ID to the resource policy.
B. In the company's AWS account create an IAM role that trusts the auditors' AWS account Create an IAM policy that has the required permissions. Attach the policy to the role. Assign a unique external ID to the role's trust policy.
C. In the company's AWS account, create an IAM user. Attach the required IAM policies to the IAM user. Create API access keys for the IAM user. Share the access keys with the auditors.
D. 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 auditor. Add the IAM users to the IAM group.

Correct Answer: B
Section:
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.
AWS IAM Roles documentation: https://aws.amazon.com/iam/features/roles/
AWS IAM Best practices: https://aws.amazon.com/iam/security-best-practices/

www.VCEplus.io

QUESTION 105
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) cluster. Configure an application to read data by using DAX and to write data directly to the DynamoDB table.
C. Create a three-node DynamoDB Accelerator (DAX) cluster. Configure an application to read data directly from the DynamoDB table and to write data by using DAX.
D. Create a single-node DynamoD8 Accelerator (DAX) cluster. Configure an application to read data by using DAX and to write data directly to the DynamoD8 table.

Correct Answer: B
Section:
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

QUESTION 106
A company has migrated an application from on premises to AWS. The application frontend is a static website that runs on two Amazon EC2 instances behind an Application Load Balancer (ALB). The application backend is a Python application that runs on three EC2 instances behind another ALB. The EC2 instances are large, general purpose On-Demand Instances that were sized to meet the on-premises specifications for

CEplus

peak usage of the application.
The application averages hundreds of thousands of requests each month. However, the application is used mainly during lunchtime and receives minimal traffic during the rest of the day.
A solutions architect needs to optimize the infrastructure cost of the application without negatively affecting the application availability.
Which combination of steps will meet these requirements? (Choose two.)

A. Change all the EC2 instances to compute optimized instances that have the same number of cores as the existing EC2 instances.
B. Move the application frontend to a static website that is hosted on Amazon S3.
C. Deploy the application frontend by using AWS Elastic Beanstalk. Use the same instance type for the nodes.
D. Change all the backend EC2 instances to Spot Instances.
E. Deploy the backend Python application to general purpose burstable EC2 instances that have the same number of cores as the existing EC2 instances.

Correct Answer: B, D
Section:
Explanation:
Moving the application frontend to a static website that is hosted on Amazon S3 will save cost as S3 is cheaper than running EC2 instances.
Using Spot instances for the backend EC2 instances will also save cost, as they are significantly cheaper than On-Demand instances. This will be suitable for the application, as it has minimal traffic during the rest of the day, and the availability of spot instances will not negatively affect the application's availability.
Amazon S3 pricing: https://aws.amazon.com/s3/pricing/
Amazon EC2 Spot Instances documentation: https://aws.amazon.com/ec2/spot/
AWS Elastic Beanstalk documentation: https://aws.amazon.com/elasticbeanstalk/
Amazon Elastic Compute Cloud (EC2) pricing: https://aws.amazon.com/ec2/pricing/

QUESTION 107
A company is running an event ticketing platform on AWS and wants to optimize the platform's cost-effectiveness. The platform is deployed on Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 and is backed by an Amazon RDS for MySQL DB instance. The company is developing new application features to run on Amazon EKS with AWS Fargate.
The platform experiences infrequent high peaks in demand. The surges in demand depend on event dates.
Which solution will provide the MOST cost-effective setup for the platform?

A. Purchase Standard Reserved Instances for the EC2 instances that the EKS cluster uses in its baseline load. Scale the cluster with Spot Instances to handle peaks. Purchase 1-year All Upfront Reserved Instances for the database to meet predicted peak load for the year.
B. Purchase Compute Savings Plans for the predicted medium load of the EKS cluster. Scale the cluster with On-Demand Capacity Reservations based on event dates for peaks. Purchase 1-year No Upfront Reserved Instances for the database to meet the predicted base load. Temporarily scale out database read replicas during peaks.
C. Purchase EC2 Instance Savings Plans for the predicted base load of the EKS cluster. Scale the cluster with Spot Instances to handle peaks. Purchase 1-year All Upfront Reserved Instances for the database to meet the predicted base load. Temporarily scale up the DB instance manually during peaks.
D. Purchase Compute Savings Plans for the predicted base load of the EKS cluster. Scale the cluster with Spot Instances to handle peaks. Purchase 1-year All Upfront Reserved Instances for the database to meet the predicted base load. Temporarily scale up the DB instance manually during peaks.

Correct Answer: B
Section:
Explanation:
They all mention using spot instances and EKS based on EC2. A spot instance is not appropriate for a production server and the company is developing new application designed for AWS Fargate, which means we must plan the future cost improvement including AWS Fargate. https://aws.amazon.com/savingsplans/compute-pricing/

QUESTION 108
A company has deployed an application on AWS Elastic Beanstalk. The application uses Amazon Aurora for the database layer. An Amazon CloudFront distribution serves web requests and includes the Elastic Beanstalk domain name as the origin server. The distribution is configured with an alternate domain name that visitors use when they access the application.
Each week, the company takes the application out of service for routine maintenance. During the time that the application is unavailable, the company wants visitors to receive an informational message instead of

VCEûp CEplus IT Certification Exams - Questions & Answers | VCEplus.io CEplus VCEûp

CEplus

a CloudFront error message.
A solutions architect creates an Amazon S3 bucket as the first step in the process.
Which combination of steps should the solutions architect take next to meet the requirements? (Choose three.)

A. Upload static informational content to the S3 bucket.
B. Create a new CloudFront distribution. Set the S3 bucket as the origin.
C. Set the S3 bucket as a second origin in the original CloudFront distribution. Configure the distribution and the S3 bucket to use an origin access identity (OAI).
D. During the weekly maintenance, edit the default cache behavior to use the S3 origin. Revert the change when the maintenance is complete.
E. During the weekly maintenance, create a cache behavior for the S3 origin on the new distribution. Set the path pattern to \ Set the precedence to 0. Delete the cache behavior when the maintenance is complete.
F. During the weekly maintenance, configure Elastic Beanstalk to serve traffic from the S3 bucket.

Correct Answer: A, C, D
Section:
Explanation:
The company wants to serve static content from an S3 bucket during the maintenance period. To do this, the following steps are required:
Upload static informational content to the S3 bucket. This will provide the source of the content that will be served to the visitors.
Set the S3 bucket as a second origin in the original CloudFront distribution. Configure the distribution and the S3 bucket to use an origin access identity (OAI). This will allow CloudFront to access the S3 bucket securely and prevent public access to the bucket.
During the weekly maintenance, edit the default cache behavior to use the S3 origin. Revert the change when the maintenance is complete. This will redirect all web requests to the S3 bucket instead of the Elastic Beanstalk domain name.
The other options are not correct because:
Creating a new CloudFront distribution is not necessary and would require changing the alternate domain name configuration.
Creating a cache behavior for the S3 origin on a new distribution would not work because the visitors would still access the original distribution using the alternate domain name.
Configuring Elastic Beanstalk to serve traffic from the S3 bucket is not possible and would not achieve the desired result.
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/DownloadDistS3AndCustomOrigins.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/distribution-web-values-specify.html#DownloadDistValuesPathPattern

QUESTION 109
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 version. Use the SLATEST version to test image processing parameters.
B. Create an Amazon DynamoDB table to store the image processing parameters. Modify the Lambda function to retrieve the image processing parameters from the DynamoDB table.
C. Directly code the image processing parameters within the Lambda function and remove the environment variables. Publish a new function version when the company updates the parameters.
D. Create a Lambda function alias. Modify the client application to use the function alias ARN. Reconfigure the Lambda alias to point to new versions of the function when the company finishes testing.

Correct Answer: D
Section:
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

VCEplus

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.
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/

QUESTION 110
A global media company is planning a multi-Region deployment of an application. Amazon DynamoDB global tables will back the deployment to keep the user experience consistent across the two continents where users are concentrated. Each deployment will have a public Application Load Balancer (ALB). The company manages public DNS internally. The company wants to make the application available through an apex domain.
Which solution will meet these requirements with the LEAST effort?

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

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

QUESTION 111
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 image. Store the image in an S3 bucket. Create a Lambda layer that uses the Docker image as the source. Deploy the API's Lambda functions as Zip packages. Configure the packages to use the Lambda layer.
B. Deploy the shared libraries and custom classes to a Docker image. Upload the image to Amazon Elastic Container Registry (Amazon ECR). Create a Lambda layer that uses the Docker image as the source. Deploy the API's Lambda functions as Zip packages. Configure the packages to use the Lambda layer.
C. Deploy the shared libraries and custom classes to a Docker container in Amazon Elastic Container Service (Amazon ECS) by using the AWS Fargate launch type. Deploy the API's Lambda functions as Zip packages. Configure the packages to use the deployed container as a Lambda layer.
D.

VCEûp VCEplus IT Certification Exams - Questions & Answers | VCEplus.io VCEplus VCEûp

CEplus

Correct Answer: B
Section:
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.
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/

QUESTION 112
A manufacturing company is building an inspection solution for its factory. The company has IP cameras at the end of each assembly line. The company has used Amazon SageMaker to train a machine learning (ML) model to identify common defects from still images.
The company wants to provide local feedback to factory workers when a defect is detected. The company must be able to provide this feedback even if the factory's internet connectivity is down. The company has a local Linux server that hosts an API that provides local feedback to the workers.
How should the company deploy the ML model to meet these requirements?

A. Set up an Amazon Kinesis video stream from each IP camera to AWS. Use Amazon EC2 instances to take still images of the streams. Upload the images to an Amazon S3 bucket. Deploy a SageMaker endpoint with the ML model. Invoke an AWS Lambda function to call the inference endpoint when new images are uploaded. Configure the Lambda function to call the local API when a defect is detected.
B. Deploy AWS IoT Greengrass on the local server. Deploy the ML model to the Greengrass server. Create a Greengrass component to take still images from the cameras and run inference. Configure the component to call the local API when a defect is detected.
C. Order an AWS Snowball device. Deploy a SageMaker endpoint the ML model and an Amazon EC2 instance on the Snowball device. Take still images from the cameras. Run inference from the EC2 instance. Configure the instance to call the local API when a defect is detected.
D. Deploy Amazon Monitron devices on each IP camera. Deploy an Amazon Monitron Gateway on premises. Deploy the ML model to the Amazon Monitron devices. Use Amazon Monitron health state alarms to call the local API from an AWS Lambda function when a defect is detected.

Correct Answer: B
Section:
Explanation:
The company should use AWS IoT Greengrass to deploy the ML model to the local server and provide local feedback to the factory workers.AWS IoT Greengrass is a service that extends AWS cloud capabilities to local devices, allowing them to collect and analyze data closer to the source of information, react autonomously to local events, and communicate securely with each other on local networks1.AWS IoT Greengrass also supports ML inference at the edge, enabling devices to run ML models locally without requiring internet connectivity2.
The other options are not correct because:
Setting up an Amazon Kinesis video stream from each IP camera to AWS would not work if the factory's internet connectivity is down. It would also incur unnecessary costs and latency to stream video data to the cloud and back.
Ordering an AWS Snowball device would not be a scalable or cost-effective solution for deploying the ML model.AWS Snowball is a service that provides physical devices for data transfer and edge computing, but it is not designed for continuous operation or frequent updates3.
Deploying Amazon Monitron devices on each IP camera would not work because Amazon Monitron is a service that monitors the condition and performance of industrial equipment using sensors and machine learning, not cameras4.
https://aws.amazon.com/greengrass/
https://docs.aws.amazon.com/greengrass/v2/developerguide/use-machine-learning-inference.html
https://aws.amazon.com/snowball/

VCEplus

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

QUESTION 113
A solutions architect needs to migrate an on-premises legacy application to AWS. The application runs on two servers behind a bad balancer. The application requires a license file that is associated with the MAC address of the server's network adapter. It takes the software vendor 12 hours to send new license files. The application also uses configuration files with a static IP address to access a database host names are not supported.
Given these requirements. which combination of steps should be taken to implement highly available architecture for the application servers in AWS? (Select TWO.)

A. Create a pool of ENIs. Request license files from the vendor for the pool, and store the license files in Amazon $3. Create a bootstrap automation script to download a license file and attach the corresponding ENI to an Amazon EC2 instance.
B. Create a pool of ENIs. Request license files from the vendor for the pool, store the license files on an Amazon EC2 instance. Create an AMI from the instance and use this AMI for all future EC2
C. Create a bootstrap automation script to request a new license file from the vendor. When the response is received, apply the license file to an Amazon EC2 instance.
D. Edit the bootstrap automation script to read the database server IP address from the AWS Systems Manager Parameter Store. and inject the value into the local configuration files.
E. Edit an Amazon EC2 instance to include the database server IP address in the configuration files and re-create the AMI to use for all future EC2 instances.

Correct Answer: A, D
Section:
Explanation:
This solution will meet the requirements of implementing a highly available architecture for the application servers in AWS. Creating a pool of ENIs will allow the application servers to have consistent MAC addresses, which are needed for the license files. Requesting license files from the vendor for the pool and storing them in Amazon S3 will ensure that the license files are available and secure. Creating a bootstrap automation script to download a license file and attach the corresponding ENI to an EC2 instance will automate the process of launching new application servers with valid licenses. Editing the bootstrap automation script to read the database server IP address from the AWS Systems Manager Parameter Store and inject the value into the local configuration files will enable the application servers to access the database without hard-coding the IP address in the configuration files. This will also allow changing the database server IP address without modifying the configuration files on each application server.

www.VCEplus.io

QUESTION 114
A company migrated an application to the AWS Cloud. The application runs on two Amazon EC2 instances behind an Application Load Balancer (ALB). Application data is stored in a MySQL database that runs on an additional EC2 instance. The application's use of the database is read-heavy.
The loads static content from Amazon Elastic Block Store (Amazon EBS) volumes that are attached to each EC2 instance. The static content is updated frequently and must be copied to each EBS volume.
The load on the application changes throughout the day. During peak hours, the application cannot handle all the incoming requests. Trace data shows that the database cannot handle the read load during peak hours.
Which solution will improve the reliability of the application?

A. Migrate the application to a set of AWS Lambda functions. Set the Lambda functions as targets for the ALB. Create a new single EBS volume for the static content. Configure the Lambda functions to read from the new EBS volume. Migrate the database to an Amazon RDS for MySQL Multi-AZ DB cluster.
B. Migrate the application to a set of AWS Step Functions state machines. Set the state machines as targets for the ALB. Create an Amazon Elastic File System (Amazon EFS) file system for the static content. Configure the state machines to read from the EFS file system. Migrate the database to Amazon Aurora MySQL Serverless v2 with a reader DB instance.
C. Containerize the application. Migrate the application to an Amazon Elastic Container Service (Amazon ECS) Cluster. Use the AWS Fargate launch type for the tasks that host the application. Create a new single EBS volume the static content. Mount the new EBS volume on the ECS duster. Configure AWS Application Auto Scaling on ECS cluster. Set the ECS service as a target for the ALB. Migrate the database to an Amazon RDS for MySOL Multi-AZ DB cluster.
D. Containerize the application. Migrate the application to an Amazon Elastic Container Service (Amazon ECS) cluster. Use the AWS Fargate launch type for the tasks that host the application. Create an Amazon Elastic File System (Amazon EFS) file system for the static content. Mount the EFS file system to each container. Configure AWS Application Auto Scaling on the ECS cluster Set the ECS service as a target for the ALB. Migrate the database to Amazon Aurora MySQL Serverless v2 with a reader DB instance.

Correct Answer: D
Section:
Explanation:
This solution will improve the reliability of the application by addressing the issues of scalability, availability, and performance. Containerizing the application will make it easier to deploy and manage on AWS.

VCEûp VCEplus IT Certification Exams - Questions & Answers | VCEplus.io VCEplus VCEûp

VCEplus

Migrating the application to an Amazon ECS cluster will allow the application to run on a fully managed container orchestration service. Using the AWS Fargate launch type for the tasks that host the application will enable the application to run on serverless compute engines that are automatically provisioned and scaled by AWS. Creating an Amazon EFS file system for the static content will provide a scalable and shared storage solution that can be accessed by multiple containers. Mounting the EFS file system to each container will eliminate the need to copy the static content to each EBS volume and ensure that the content is always up to date. Configuring AWS Application Auto Scaling on the ECS cluster will enable the application to scale up and down based on demand or a predefined schedule. Setting the ECS service as a target for the ALB will distribute the incoming requests across multiple tasks in the ECS cluster and improve the availability and fault tolerance of the application. Migrating the database to Amazon Aurora MySQL Serverless v2 with a reader DB instance will provide a fully managed, compatible, and scalable relational database service that can handle high throughput and concurrent connections. Using a reader DB instance will offload some of the read load from the primary DB instance and improve the performance of the database.

QUESTION 115
A company has AWS accounts that are in an organization in AWS Organizations. The company wants to track Amazon EC2 usage as a metric. The company's architecture team must receive a daily alert if the EC2 usage is more than 10% higher than the average EC2 usage from the last 30 days.
Which solution will meet these requirements?

A. Configure AWS Budgets in the organization's management account. Specify a usage type of EC2 running hours. Specify a daily period. Set the budget amount to be 10% more than the reported average usage for the last 30 days from AWS Cost Explorer. Configure an alert to notify the architecture team if the usage threshold is met.
B. Configure AWS Cost Anomaly Detection in the organization's management account. Configure a monitor type of AWS Service. Apply a filter of Amazon EC2. Configure an alert subscription to notify the architecture team if the usage is 10% more than the average usage for the last 30 days.
C. Enable AWS Trusted Advisor in the organization's management account. Configure a cost optimization advisory alert to notify the architecture team if the EC2 usage is 10% more than the reported average usage for the last 30 days.
D. Configure Amazon Detective in the organization's management account. Configure an EC2 usage anomaly alert to notify the architecture team if Detective identifies a usage anomaly of more than 10%.

Correct Answer: B
Section:
Explanation:
AWS Cost Anomaly Detection is a feature of the AWS Cost Management suite that leverages machine learning to enable continuous monitoring of your AWS costs and usage, allowing you to identify unexpected and abnormal spending1.You can create cost monitors that evaluate specific AWS services, member accounts, cost allocation tags, or cost categories based on your AWS account structure2.You can also configure alert subscriptions that notify you when a cost monitor detects an anomaly that meets your threshold2. In this case, you can create a cost monitor with a monitor type of AWS Service and apply a filter of Amazon EC2 to track the EC2 usage as a metric.You can then configure an alert subscription to notify the architecture team if the usage is 10% more than the average usage for the last 30 days, which is the anomaly detection period used by AWS Cost Anomaly Detection3.

QUESTION 116
A company is running an application on premises. The application uses a set of web servers that host a static React-based single-page application (SPA), a Node.js API, and a MYSQL database server. The database is read intensive. The company will need to expand the database's storage at an unpredictable rate.
The company must migrate the application to AWS. The company also must modernize the architecture to reduce infrastructure management and increase scalability.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use AWS Database Migration Service (AWS DMS) to migrate the database to Amazon RDS for MySQL. Use AWS Application Migration Service to migrate the web application to a fleet of Amazon EC2 instances behind an Elastic Load Balancing (ELB) load balancer. Use a Spot Fleet with a request type of request to host the API.
B. Use AWS Database Migration Service (AWS DMS) to migrate the database to Amazon Aurora MySQL. Copy the web files to an Amazon S3 bucket and set up web hosting. Copy the API code to AWS Lambda functions. Configure Amazon API Gateway to point to the Lambda functions.
C. Use AWS Database Migration Service (AWS DMS) to migrate the database to a MySQL database that runs on Amazon EC2 instances. Use AWS DataSync to migrate the web files and API files to an Amazon FSx for Windows File Server file system. Set up a fleet of EC2 instances in an Auto Scaling group as web servers. Mount the FSx for Windows File Server file system.
D. Use AWS Application Migration Service to migrate the database to Amazon EC2 instances. Copy the web files to containers that run on Amazon Elastic Kubernetes Service (Amazon EKS). Set up an Elastic Load Balancing (ELB) load balancer for the EC2 instances and EKS containers. Copy the API code to AWS Lambda functions. Configure Amazon API Gateway to point to the Lambda functions.

Correct Answer: B
Section:

VCEûp VCEplus IT Certification Exams - Questions & Answers | VCEplus.io VCEplus VCEûp

VCEplus

QUESTION 117
A company is using AWS Control Tower to manage AWS accounts in an organization in AWS Organizations. The company has an OU that contains accounts. The company must prevent any new or existing Amazon EC2 instances in the OUs accounts from gaining a public IP address.
Which solution will meet these requirements?

A. Configure all instances in each account in the OU to use AWS Systems Manager. Use a Systems Manager Automation runbook to prevent public IP addresses from being attached to the instances.
B. Implement the AWS Control Tower proactive control to check whether instances in the OU's accounts have a public IP address. Set the AssociatePublicIpAddress property to False. Attach the proactive control to the OU.
C. Create an SCP that prevents the launch of instances that have a public IP address. Additionally, configure the SCP to prevent the attachment of a public IP address to existing instances. Attach the SCP to the OU.
D. Create an AWS Config custom rule that detects instances that have a public IP address. Configure a remediation action that uses an AWS Lambda function to detach the public IP addresses from the instances.

Correct Answer: C
Section:
Explanation:
This option will meet the requirements of preventing any new or existing EC2 instances in the OU's accounts from gaining a public IP address. An SCP is a policy that you can attach to an OU or an account in AWS Organizations to define the maximum permissions for the entities in that OU or account. By creating an SCP that denies the ec2:RunInstances and ec2:AssociateAddress actions when the value of the aws:RequestTag/aws:PublicIp condition key is true, you can prevent any user or role in the OU from launching instances that have a public IP address or attaching a public IP address to existing instances. This will effectively enforce a security best practice and reduce the risk of unauthorized access to your EC2 instances.

www.VCEplus.io