



Amazon-Web-Services

Exam Questions SAP-C02

AWS Certified Solutions Architect - Professional

NEW QUESTION 1

- (Exam Topic 1)

A startup company recently migrated a large ecommerce website to AWS. The website has experienced a 70% increase in sales. Software engineers are using a private GitHub repository to manage code. The DevOps learn is using Jenkins for builds and unit testing. The engineers need to receive notifications for bad builds and zero downtime during deployments. The engineers also need to ensure any changes to production are seamless for users and can be rolled back in the event of a major issue.

The software engineers have decided to use AWS CodePipeline to manage their build and deployment process.

Which solution will meet these requirements?

- A. Use GitHub websockets to trigger the CodePipeline pipeline
- B. Use the Jenkins plugin for AWS CodeBuild to conduct unit testing
- C. Send alerts to an Amazon SNS topic for any bad build
- D. Deploy in an in-place
- E. all-at-once deployment configuration using AWS CodeDeploy.
- F. Use GitHub webhooks to trigger the CodePipeline pipeline
- G. Use the Jenkins plugin for AWS CodeBuild to conduct unit testing
- H. Send alerts to an Amazon SNS topic for any bad build
- I. Deploy in a blue/green deployment using AWS CodeDeploy.
- J. Use GitHub websockets to trigger the CodePipeline pipeline
- K. Use AWS X-Ray for unit testing and static code analysis
- L. Send alerts to an Amazon SNS topic for any bad build
- M. Deploy in a blue/green deployment using AWS CodeDeploy.
- N. Use GitHub webhooks to trigger the CodePipeline pipeline
- O. Use AWS X-Ray for unit testing and static code analysis
- P. Send alerts to an Amazon SNS topic for any bad build
- Q. Deploy in an in-place, all-at-once deployment configuration using AWS CodeDeploy.

Answer: B

NEW QUESTION 2

- (Exam Topic 1)

A company is serving files to its customers through an SFTP server that is accessible over the internet. The SFTP server is running on a single Amazon EC2 instance with an Elastic IP address attached. Customers connect to the SFTP server through its Elastic IP address and use SSH for authentication. The EC2 instance also has an attached security group that allows access from all customer IP addresses.

A solutions architect must implement a solution to improve availability, minimize the complexity of infrastructure management, and minimize the disruption to customers who access files. The solution must not change the way customers connect.

Which solution will meet these requirements?

- A. Disassociate the Elastic IP address from the EC2 instance. Create an Amazon S3 bucket to be used for SFTP file hosting. Create an AWS Transfer Family server. Configure the Transfer Family server with a publicly accessible endpoint.
- B. Associate the SFTP Elastic IP address with the new endpoint.
- C. Point the Transfer Family server to the S3 bucket. Sync all files from the SFTP server to the S3 bucket.
- D. Disassociate the Elastic IP address from the EC2 instance.
- E. Create an Amazon S3 bucket to be used for SFTP file hosting. Create an AWS Transfer Family server.
- F. Configure the Transfer Family server with a VPC-hosted endpoint.
- G. Internet-facing endpoint.
- H. Associate the SFTP Elastic IP address with the new endpoint.
- I. Attach the security group with customer IP addresses to the new endpoint.
- J. Point the Transfer Family server to the S3 bucket.
- K. Sync all files from the SFTP server to the S3 bucket.
- L. Disassociate the Elastic IP address from the EC2 instance.
- M. Create a new Amazon Elastic File System (Amazon EFS) file system to be used for SFTP file hosting.
- N. Create an AWS Fargate task definition to run an SFTP server.
- O. Specify the EFS file system as a mount in the task definition. Create a Fargate service by using the task definition, and place a Network Load Balancer (NLB) in front of the service. When configuring the service, attach the security group with customer IP addresses to the tasks that run the SFTP server. Associate the Elastic IP address with the NLB. Sync all files from the SFTP server to the S3 bucket.
- P. Disassociate the Elastic IP address from the EC2 instance. Create a multi-attach Amazon Elastic Block Store (Amazon EBS) volume to be used for SFTP file hosting. Create a Network Load Balancer (NLB) with the Elastic IP address attached. Create an Auto Scaling group with EC2 instances that run an SFTP server. Define in the Auto Scaling group that instances that are launched should attach the new multi-attach EBS volume. Configure the Auto Scaling group to automatically add instances behind the NLB. Configure the Auto Scaling group to use the security group that allows customer IP addresses for the EC2 instances that the Auto Scaling group launches. Sync all files from the SFTP server to the new multi-attach EBS volume.

Answer: B

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/aws-sftp-endpoint-type/>

<https://docs.aws.amazon.com/transfer/latest/userguide/create-server-in-vpc.html> <https://aws.amazon.com/premiumsupport/knowledge-center/aws-sftp-endpoint-type/>

NEW QUESTION 3

- (Exam Topic 1)

A financial services company logs personally identifiable information in its application logs stored in Amazon S3. Due to regulatory compliance requirements, the log files must be encrypted at rest. The security team has mandated that the company's on-premises hardware security modules (HSMs) be used to generate the CMK material.

Which steps should the solutions architect take to meet these requirements?

- A. Create an AWS CloudHSM cluster
- B. Create a new CMK in AWS KMS using AWS_CloudHSM as the source (or the key material and an origin of AWS_CLOUDHSM)
- C. Enable automatic key rotation on the CMK with a duration of 1 year

- D. Configure a bucket policy on the logging bucket that disallows uploads of unencrypted data and requires that the encryption source be AWS KMS.
- E. Provision an AWS Direct Connect connection, ensuring there is no overlap of the RFC 1918 address space between on-premises hardware and the VPC
- F. Configure an AWS bucket policy on the logging bucket that requires all objects to be encrypted
- G. Configure the logging application to query on-premises HSMs from the AWS environment for the encryption key material, and create a unique CMK for each logging event.
- H. Create a CMK in AWS KMS with no key material and an origin of EXTERNAL
- I. Import the key material generated from the on-premises HSMs into the CMK using the public key and import token provided by AWS
- J. Configure a bucket policy on the logging bucket that disallows uploads of non-encrypted data and requires that the encryption source be AWS KMS.
- K. Create a new CMK in AWS KMS with AWS-provided key material and an origin of AWS_KM
- L. Disable this CM
- M. and overwrite the key material with the key material from the on-premises HSM using the public key and import token provided by AWS
- N. Re-enable the CM
- O. Enable automatic key rotation on the CMK with a duration of 1 year
- P. Configure a bucket policy on the logging bucket that disallows uploads of non-encrypted data and requires that the encryption source be AWS KMS.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/security/how-to-byok-bring-your-own-key-to-aws-kms-for-less-than-15-00-a-year/>
<https://docs.aws.amazon.com/kms/latest/developerguide/importing-keys-create-cmk.html>

NEW QUESTION 4

- (Exam Topic 1)

An enterprise runs 103 line-of-business applications on virtual machines in an on-premises data center. Many of the applications are simple PHP, Java, or Ruby web applications, are no longer actively developed, and serve little traffic.

Which approach should be used to migrate these applications to AWS with the LOWEST infrastructure costs?

- A. Deploy the applications to single-instance AWS Elastic Beanstalk environments without a load balancer.
- B. Use AWS SMS to create AMIs for each virtual machine and run them in Amazon EC2.
- C. Convert each application to a Docker image and deploy to a small Amazon ECS cluster behind an Application Load Balancer.
- D. Use VM Import/Export to create AMIs for each virtual machine and run them in single-instance AWS Elastic Beanstalk environments by configuring a custom image.

Answer: C

NEW QUESTION 5

- (Exam Topic 1)

A company has an application that sells tickets online and experiences bursts of demand every 7 days. The application has a stateless presentation layer running on Amazon EC2, an Oracle database to store unstructured data catalog information, and a backend API layer. The front-end layer uses an Elastic Load Balancer to distribute the load across nine On-Demand Instances over three Availability Zones (AZs). The Oracle database is running on a single EC2 instance. The company is experiencing performance issues when running more than two concurrent campaigns. A solutions architect must design a solution that meets the following requirements:

- Address scalability issues.
- Increase the level of concurrency.
- Eliminate licensing costs.
- Improve reliability.

Which set of steps should the solutions architect take?

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

Answer: C

Explanation:

Combination of On-Demand and Spot Instances + DynamoDB.

NEW QUESTION 6

- (Exam Topic 1)

A company wants to deploy an AWS WAF solution to manage AWS WAF rules across multiple AWS accounts. The accounts are managed under different OUs in AWS Organizations.

Administrators must be able to add or remove accounts or OUs from managed AWS WAF rule sets as needed. Administrators also must have the ability to automatically update and remediate noncompliant AWS WAF rules in all accounts.

Which solution meets these requirements with the LEAST amount of operational overhead?

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

F. Use AWS Control Tower to manage AWS WAF rules across accounts in the organization Use AWS Key Management Service (AWS KMS) to store account numbers and OUs to manage Update AWS KMS as needed to add or remove accounts or OUs Create IAM users in member accounts Allow AWS Control Tower in the management account to use the access key and secret access key to create and update AWS WAF rules in the member accounts

Answer: B

NEW QUESTION 7

- (Exam Topic 1)

A company wants to retire its Oracle Solaris NFS storage arrays. The company requires rapid data migration over its internet network connection to a combination of destinations for Amazon S3, Amazon Elastic File System (Amazon EFS), and Amazon FSx for Windows File Server. The company also requires a full initial copy, as well as incremental transfers of changes until the retirement of the storage arrays. All data must be encrypted and checked for integrity. What should a solutions architect recommend to meet these requirements?

- A. Configure CloudEndur
- B. Create a project and deploy the CloudEndure agent and token to the storage arra
- C. Run the migration plan to start the transfer.
- D. Configure AWS DataSyn
- E. Configure the DataSync agent and deploy it to the local networ
- F. Create a transfer task and start the transfer.
- G. Configure the aws S3 sync comman
- H. Configure the AWS client on the client side with credential
- I. Run the sync command to start the transfer.
- J. Configure AWS Transfer (or FT
- K. Configure the FTP client with credential
- L. Script the client to connect and sync to start the transfer.

Answer: B

NEW QUESTION 8

- (Exam Topic 1)

A company has an Amazon VPC that is divided into a public subnet and a private subnet. A web application runs in Amazon VPC, and each subnet has its own NACL. The public subnet has a CIDR of 10.0.0.0/24. An Application Load Balancer is deployed to the public subnet. The private subnet has a CIDR of 10.0.1.0/24. Amazon EC2 instances that run a web server on port 80 are launched into the private subnet. Only network traffic that is required for the Application Load Balancer to access the web application can be allowed to travel between the public and private subnets. What collection of rules should be written to ensure that the private subnet's NACL meets the requirement? (Select TWO.)

- A. An inbound rule for port 80 from source 0.0.0.0/0
- B. An inbound rule for port 80 from source 10.0.0.0/24
- C. An outbound rule for port 80 to destination 0.0.0.0/0
- D. An outbound rule for port 80 to destination 10.0.0.0/24
- E. An outbound rule for ports 1024 through 65535 to destination 10.0.0.0/24

Answer: BE

Explanation:

Ephemeral ports are not covered in the syllabus so be careful that you don't confuse day to day best practise with what is required for the exam. Link to an explanation on Ephemeral ports here. <https://acloud.guru/forums/aws-certified-solutions-architect-associate/discussion/-KUbcwo4IXefMI7janaK/netw>

NEW QUESTION 9

- (Exam Topic 1)

A travel company built a web application that uses Amazon Simple Email Service (Amazon SES) to send email notifications to users. The company needs to enable logging to help troubleshoot email delivery issues. The company also needs the ability to do searches that are based on recipient, subject, and time sent. Which combination of steps should a solutions architect take to meet these requirements? (Select TWO.)

- A. Create an Amazon SES configuration set with Amazon Kinesis Data Firehose as the destination
- B. Choose to send logs to an Amazon S3 bucket.
- C. Enable AWS CloudTrail logging
- D. Specify an Amazon S3 bucket as the destination for the logs.
- E. Use Amazon Athena to query the logs in the Amazon S3 bucket for recipient, subject, and time sent.
- F. Create an Amazon CloudWatch log group
- G. Configure Amazon SES to send logs to the log group
- H. Use Amazon Athena to query the logs in Amazon CloudWatch for recipient, subject, and time sent.

Answer: AC

Explanation:

<https://docs.aws.amazon.com/ses/latest/dg/event-publishing-retrieving-firehose.html>

To enable you to track your email sending at a granular level, you can set up Amazon SES to publish email sending events to Amazon CloudWatch, Amazon Kinesis Data Firehose, or Amazon Simple Notification Service based on characteristics that you define.

<https://docs.aws.amazon.com/ses/latest/dg/monitor-using-event-publishing.html>

<https://aws.amazon.com/getting-started/hands-on/build-serverless-real-time-data-processing-app-lambda-kinesis>

NEW QUESTION 10

- (Exam Topic 1)

A company has a three-tier application running on AWS with a web server, an application server, and an Amazon RDS MySQL DB instance. A solutions architect is designing a disaster recovery (DR) solution with an RPO of 5 minutes. Which solution will meet the company's requirements?

- A. Configure AWS Backup to perform cross-Region backups of all servers every 5 minute
- B. Reprovision the three tiers in the DR Region from the backups using AWS CloudFormation in the event of a disaster.
- C. Maintain another running copy of the web and application server stack in the DR Region using AWS CloudFormation drill detectio
- D. Configure cross-Region snapshots of the DB instance to the DR Region every 5 minute
- E. In the event of a disaster, restore the DB instance using the snapshot in the DR Region.
- F. Use Amazon EC2 Image Builder to create and copy AMIs of the web and application server to both the primary and DR Region
- G. Create a cross-Region read replica of the DB instance in the DR Region
- H. In the event of a disaster, promote the read replica to become the master and reprovision the servers with AWS CloudFormation using the AMIs.
- I. Create AMIs of the web and application servers in the DR Region
- J. Use scheduled AWS Glue jobs to synchronize the DB instance with another DB instance in the DR Region
- K. In the event of a disaster, switch to the DB instance in the DR Region and reprovision the servers with AWS CloudFormation using the AMIs.

Answer: C

Explanation:

deploying a brand new RDS instance will take >30 minutes. You will use EC2 Image builder to put the AMIs into the new region, but not use image builder to LAUNCH them.

NEW QUESTION 10

- (Exam Topic 1)

A company has multiple AWS accounts as part of an organization created with AWS Organizations. Each account has a VPC in the us-east-2 Region and is used for either production or development workloads. Amazon EC2 instances across production accounts need to communicate with each other, and EC2 instances across development accounts need to communicate with each other, but production and development instances should not be able to communicate with each other.

To facilitate connectivity, the company created a common network account. The company used AWS Transit Gateway to create a transit gateway in the us-east-2 Region in the network account and shared the transit gateway with the entire organization by using AWS Resource Access Manager. Network administrators then attached VPCs in each account to the transit gateway, after which the EC2 instances were able to communicate across accounts. However, production and development accounts were also able to communicate with one another.

Which set of steps should a solutions architect take to ensure production traffic and development traffic are completely isolated?

- A. Modify the security groups assigned to development EC2 instances to block traffic from production EC2 instance
- B. Modify the security groups assigned to production EC2 instances to block traffic from development EC2 instances.
- C. Create a tag on each VPC attachment with a value of either production or development, according to the type of account being attached
- D. Using the Network Manager feature of AWS Transit Gateway, create policies that restrict traffic between VPCs based on the value of this tag.
- E. Create separate route tables for production and development traffic
- F. Delete each account's association and route propagation to the default AWS Transit Gateway route table
- G. Attach development VPCs to the development AWS Transit Gateway route table and production VPCs to the production route table, and enable automatic route propagation on each attachment.
- H. Create a tag on each VPC attachment with a value of either production or development, according to the type of account being attached
- I. Modify the AWS Transit Gateway routing table to route production tagged attachments to one another and development tagged attachments to one another.

Answer: C

Explanation:

<https://docs.aws.amazon.com/vpc/latest/tgw/vpc-tgw.pdf>

NEW QUESTION 13

- (Exam Topic 1)

A company is deploying a new cluster for big data analytics on AWS. The cluster will run across many Linux Amazon EC2 instances that are spread across multiple Availability Zones.

All of the nodes in the cluster must have read and write access to common underlying file storage. The file storage must be highly available, must be resilient, must be compatible with the Portable Operating System Interface (POSIX), and must accommodate high levels of throughput.

Which storage solution will meet these requirements?

- A. Provision an AWS Storage Gateway file gateway NFS file share that is attached to an Amazon S3 bucket
- B. Mount the NFS file share on each EC2 instance in the cluster.
- C. Provision a new Amazon Elastic File System (Amazon EFS) file system that uses General Purpose performance mode
- D. Mount the EFS file system on each EC2 instance in the cluster.
- E. Provision a new Amazon Elastic Block Store (Amazon EBS) volume that uses the io2 volume type. Attach the EBS volume to all of the EC2 instances in the cluster.
- F. Provision a new Amazon Elastic File System (Amazon EFS) file system that uses Max I/O performance mode
- G. Mount the EFS file system on each EC2 instance in the cluster.

Answer: D

NEW QUESTION 16

- (Exam Topic 1)

A company is running a web application on Amazon EC2 instances in a production AWS account. The company requires all logs generated from the web application to be copied to a central AWS account (for analysis and archiving). The company's AWS accounts are currently managed independently. Logging agents are configured on the EC2 instances to upload the log files to an Amazon S3 bucket in the central AWS account.

A solutions architect needs to provide access for a solution that will allow the production account to store log files in the central account. The central account also needs to have read access to the log files.

What should the solutions architect do to meet these requirements?

- A. Create a cross-account role in the central account
- B. Assume the role from the production account when the logs are being copied.
- C. Create a policy on the S3 bucket with the production account ID as the principal
- D. Allow S3 access from a delegated user.
- E. Create a policy on the S3 bucket with access from only the CIDR range of the EC2 instances in the production account
- F. Use the production account ID as the principal.

- G. Create a cross-account role in the production account
- H. Assume the role from the production account when the logs are being copied.

Answer: B

NEW QUESTION 17

- (Exam Topic 1)

A company that tracks medical devices in hospitals wants to migrate its existing storage solution to the AWS Cloud. The company equips all of its devices with sensors that collect location and usage information. This sensor data is sent in unpredictable patterns with large spikes. The data is stored in a MySQL database running on premises at each hospital. The company wants the cloud storage solution to scale with usage.

The company's analytics team uses the sensor data to calculate usage by device type and hospital. The team needs to keep analysis tools running locally while fetching data from the cloud. The team also needs to use existing Java application and SQL queries with as few changes as possible.

How should a solutions architect meet these requirements while ensuring the sensor data is secure?

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

Answer: C

Explanation:

<https://aws.amazon.com/blogs/aws/new-data-api-for-amazon-aurora-serverless/> <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/data-api.html>
<https://aws.amazon.com/blogs/aws/aws-privatelink-for-amazon-s3-now-available/> <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/data-api.html#data-api.access>

The data is currently stored in a MySQL database running on-prem. Storing MySQL data in S3 doesn't sound good so B & D are out. Aurora Data API "enables the SQL HTTP endpoint, a connectionless Web Service API for running SQL queries against this database. When the SQL HTTP endpoint is enabled, you can also query your database from inside the RDS console (these features are free to use)."

NEW QUESTION 18

- (Exam Topic 1)

A large company with hundreds of AWS accounts has a newly established centralized internal process for purchasing new or modifying existing Reserved Instances. This process requires all business units that want to purchase or modify Reserved Instances to submit requests to a dedicated team for procurement or execution. Previously, business units would directly purchase or modify Reserved Instances in their own respective AWS accounts autonomously.

Which combination of steps should be taken to proactively enforce the new process in the MOST secure way possible? (Select TWO.)

- A. Ensure all AWS accounts are part of an AWS Organizations structure operating in all features mode.
- B. Use AWS Config to report on the attachment of an IAM policy that denies access to the ec2:PurchaseReservedInstancesOffering and ec2:ModifyReservedInstances actions.
- C. In each AWS account, create an IAM policy with a DENY rule to the ec2:PurchaseReservedInstancesOffering and ec2:ModifyReservedInstances actions.
- D. Create an SCP that contains a deny rule to the ec2:PurchaseReservedInstancesOffering and ec2:ModifyReservedInstances action
- E. Attach the SCP to each organizational unit (OU) of the AWS Organizations structure.
- F. Ensure that all AWS accounts are part of an AWS Organizations structure operating in consolidated billing features mode.

Answer: AD

Explanation:

https://docs.aws.amazon.com/organizations/latest/APIReference/API_EnableAllFeatures.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp-strategies.html

NEW QUESTION 23

- (Exam Topic 1)

A company has a new application that needs to run on five Amazon EC2 instances in a single AWS Region. The application requires high-throughput, low-latency network connections between all of the EC2 instances where the application will run. There is no requirement for the application to be fault tolerant.

Which solution will meet these requirements?

- A. Launch five new EC2 instances into a cluster placement group
- B. Ensure that the EC2 instance type supports enhanced networking.
- C. Launch five new EC2 instances into an Auto Scaling group in the same Availability Zone
- D. Attach an extra elastic network interface to each EC2 instance.
- E. Launch five new EC2 instances into a partition placement group
- F. Ensure that the EC2 instance type supports enhanced networking.
- G. Launch five new EC2 instances into a spread placement group
- H. Attach an extra elastic network interface to each EC2 instance.

Answer: A

Explanation:

When you launch EC2 instances in a cluster they benefit from performance and low latency. No redundancy though as per the question
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>.

NEW QUESTION 24

- (Exam Topic 1)

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

Which solution will provide the LARGEST overall cost reduction while meeting these requirements?

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

Answer: B

NEW QUESTION 26

- (Exam Topic 1)

An e-commerce company is revamping its IT infrastructure and is planning to use AWS services. The company's CIO has asked a solutions architect to design a simple, highly available, and loosely coupled order processing application. The application is responsible (or receiving and processing orders before storing them in an Amazon DynamoDB table. The application has a sporadic traffic pattern and should be able to scale during marketing campaigns to process the orders with minimal delays.

Which of the following is the MOST reliable approach to meet the requirements?

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

Answer: B

Explanation:

Q: How does Amazon Kinesis Data Streams differ from Amazon SQS?

Amazon Kinesis Data Streams enables real-time processing of streaming big data. It provides ordering of records, as well as the ability to read and/or replay records in the same order to multiple Amazon Kinesis Applications. The Amazon Kinesis Client Library (KCL) delivers all records for a given partition key to the same record processor, making it easier to build multiple applications reading from the same Amazon Kinesis data stream (for example, to perform counting, aggregation, and filtering).

<https://aws.amazon.com/kinesis/data-streams/faqs/>

<https://aws.amazon.com/blogs/big-data/unite-real-time-and-batch-analytics-using-the-big-data-lambda-architect>

NEW QUESTION 31

- (Exam Topic 1)

A company has application services that have been containerized and deployed on multiple Amazon EC2 instances with public IPs. An Apache Kafka cluster has been deployed to the EC2 instances. A PostgreSQL database has been migrated to Amazon RDS for PostgreSQL. The company expects a significant increase of orders on its platform when a new version of its flagship product is released.

What changes to the current architecture will reduce operational overhead and support the product release?

- A. Create an EC2 Auto Scaling group behind an Application Load Balance
- B. Create additional read replicas for the DB instance
- C. Create Amazon Kinesis data streams and configure the application services to use the data stream
- D. Store and serve static content directly from Amazon S3.
- E. Create an EC2 Auto Scaling group behind an Application Load Balance
- F. Deploy the DB instance in Multi-AZ mode and enable storage auto scalin
- G. Create Amazon Kinesis data streams and configure the application services to use the data stream
- H. Store and serve static content directly from Amazon S3.
- I. Deploy the application on a Kubernetes cluster created on the EC2 instances behind an Application Load Balance
- J. Deploy the DB instance in Multi-AZ mode and enable storage auto scalin
- K. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluste
- L. Store static content in Amazon S3 behind an Amazon CloudFront distribution.
- M. Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate and enable auto scaling behind an Application Load Balance
- N. Create additional read replicas for the DB instance
- O. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluste
- P. Store static content in Amazon S3 behind an Amazon CloudFront distribution.

Answer: D

Explanation:

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

NEW QUESTION 34

- (Exam Topic 1)

A company standardized its method of deploying applications to AWS using AWS CodePipeline and AWS Cloud Formation. The applications are in Typescript and Python. The company has recently acquired another business that deploys applications to AWS using Python scripts.

Developers from the newly acquired company are hesitant to move their applications under CloudFormation because it would require than they learn a new domain-specific language and eliminate their access to language features, such as looping.

How can the acquired applications quickly be brought up to deployment standards while addressing the developers' concerns?

- A. Create CloudFormation templates and re-use parts of the Python scripts as instance user dat
- B. Use the AWS Cloud Development Kit (AWS CDK) to deploy the application using these template
- C. Incorporate the AWS CDK into CodePipeline and deploy the application to AWS using these templates.
- D. Use a third-party resource provisioning engine inside AWS CodeBuild to standardize the deployment processes of the existing and acquired compan
- E. Orchestrate the CodeBuild job using CodePipeline.
- F. Standardize on AWS OpsWork
- G. Integrate OpsWorks with CodePipelin
- H. Have the developers create Chef recipes to deploy their applications on AWS.
- I. Define the AWS resources using Typescript or Pytho
- J. Use the AWS Cloud Development Kit (AWS CDK) to create CloudFormation templates from the developers' code, and use the AWS CDK to create CloudFormation stack
- K. Incorporate the AWS CDK as a CodeBuild job in CodePipeline.

Answer: D

NEW QUESTION 35

- (Exam Topic 1)

A solutions architect is designing an application to accept timesheet entries from employees on their mobile devices. Timesheets will be submitted weekly, with most of the submissions occurring on Friday. The data must be stored in a format that allows payroll administrators to run monthly reports. The infrastructure must be highly available and scale to match the rate of incoming data and reporting requests.

Which combination of steps meets these requirements while minimizing operational overhead? (Select TWO.)

- A. Deploy the application to Amazon EC2 On-Demand Instances With load balancing across multiple Availability Zone
- B. Use scheduled Amazon EC2 Auto Scaling to add capacity before the high volume of submissions on Fridays.
- C. Deploy the application in a container using Amazon Elastic Container Service (Amazon ECS) with load balancing across multiple Availability Zone
- D. Use scheduled Service Auto Scaling to add capacity before the high volume of submissions on Fridays.
- E. Deploy the application front end to an Amazon S3 bucket served by Amazon CloudFron
- F. Deploy the application backend using Amazon API Gateway with an AWS Lambda proxy integration.
- G. Store the timesheet submission data in Amazon Redshif
- H. Use Amazon QuickSight to generate the reports using Amazon Redshift as the data source.
- I. Store the timesheet submission data in Amazon S3. Use Amazon Athena and Amazon QuickSight to generate the reports using Amazon S3 as the data source.

Answer: AE

NEW QUESTION 36

- (Exam Topic 1)

A company is developing and hosting several projects in the AWS Cloud. The projects are developed across multiple AWS accounts under the same organization in AWS Organizations. The company requires the cost lor cloud infrastructure to be allocated to the owning project. The team responsible for all of the AWS accounts has discovered that several Amazon EC2 instances are lacking the Project tag used for cost allocation.

Which actions should a solutions architect take to resolve the problem and prevent it from happening in the future? (Select THREE.)

- A. Create an AWS Config rule in each account to find resources with missing tags.
- B. Create an SCP in the organization with a deny action for ec2:RunInstances if the Project tag is missing.
- C. Use Amazon Inspector in the organization to find resources with missing tags.
- D. Create an IAM policy in each account with a deny action for ec2:RunInstances if the Project tag is missing.
- E. Create an AWS Config aggregator for the organization to collect a list of EC2 instances with the missing Project tag.
- F. Use AWS Security Hub to aggregate a list of EC2 instances with the missing Project tag.

Answer: BDE

NEW QUESTION 41

- (Exam Topic 1)

A company built an ecommerce website on AWS using a three-tier web architecture. The application is Java-based and composed of an Amazon CloudFront distribution, an Apache web server layer of Amazon EC2 instances in an Auto Scaling group, and a backend Amazon Aurora MySQL database.

Last month, during a promotional sales event, users reported errors and timeouts while adding items to their shopping carts. The operations team recovered the logs created by the web servers and reviewed Aurora DB cluster performance metrics. Some of the web servers were terminated before logs could be collected and the Aurora metrics were not sufficient for query performance analysis.

Which combination of steps must the solutions architect take to improve application performance visibility during peak traffic events? (Select THREE.)

- A. Configure the Aurora MySQL DB cluster to publish slow query and error logs to Amazon CloudWatch Logs.
- B. Implement the AWS X-Ray SDK to trace incoming HTTP requests on the EC2 instances and implement tracing of SQL queries with the X-Ray SDK for Java.
- C. Configure the Aurora MySQL DB cluster to stream slow query and error logs to Amazon Kinesis.
- D. Install and configure an Amazon CloudWatch Logs agent on the EC2 instances to send the Apache logsto CloudWatch Logs.
- E. Enable and configure AWS CloudTrail to collect and analyze application activity from Amazon EC2 and Aurora.
- F. Enable Aurora MySQL DB cluster performance benchmarking and publish the stream to AWS X-Ray.

Answer: ABD

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_LogAccess.Concepts.MySQL.html# <https://aws.amazon.com/blogs/mt/simplifying->

[apache-server-logs-with-amazon-cloudwatch-logs-insights/ https://docs.aws.amazon.com/xray/latest/devguide/xray-sdk-dotnet-messagehandler.html](https://docs.aws.amazon.com/xray/latest/devguide/xray-sdk-dotnet-messagehandler.html)
<https://docs.aws.amazon.com/xray/latest/devguide/xray-sdk-java-sqlclients.html>

NEW QUESTION 45

- (Exam Topic 1)

A large company in Europe plans to migrate its applications to the AWS Cloud. The company uses multiple AWS accounts for various business groups. A data privacy law requires the company to restrict developers' access to AWS European Regions only.

What should the solutions architect do to meet this requirement with the LEAST amount of management overhead?

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

Answer: B

Explanation:

"This policy uses the Deny effect to deny access to all requests for operations that don't target one of the two approved regions (eu-central-1 and eu-west-1)."

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_general.htm

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

NEW QUESTION 50

- (Exam Topic 1)

A company is building a hybrid solution between its existing on-premises systems and a new backend in AWS. The company has a management application to monitor the state of its current IT infrastructure and automate responses to issues. The company wants to incorporate the status of its consumed AWS services into the application. The application uses an HTTPS endpoint to receive updates.

Which approach meets these requirements with the LEAST amount of operational overhead?

- A. Configure AWS Systems Manager OpsCenter to ingest operational events from the on-premises systems Retire the on-premises management application and adopt OpsCenter as the hub
- B. Configure Amazon EventBridge (Amazon CloudWatch Events) to detect and react to changes for AWS Health events from the AWS Personal Health Dashboard Configure the EventBridge (CloudWatch Events) event to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic and subscribe the topic to the HTTPS endpoint of the management application
- C. Modify the on-premises management application to call the AWS Health API to poll for status events of AWS services.
- D. Configure Amazon EventBridge (Amazon CloudWatch Events) to detect and react to changes for AWS Health events from the AWS Service Health Dashboard Configure the EventBridge (CloudWatch Events) event to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic and subscribe the topic to an HTTPS endpoint for the management application with a topic filter corresponding to the services being used

Answer: A

Explanation:

ALB & NLB both supports IPs as targets. Questions is based on TCP traffic over VPN to on-premise. TCP is layer 4 and the , load balancer should be NLB. Then next questions does NLB supports loadbalancing traffic over VPN. And answer is YES based on below URL.

<https://aws.amazon.com/about-aws/whats-new/2018/09/network-load-balancer-now-supports-aws-vpn/>

Target as IPs for NLB & ALB: <https://aws.amazon.com/elasticloadbalancing/faqs/?nc=sn&loc=5> <https://aws.amazon.com/elasticloadbalancing/application-load-balancer/>

NEW QUESTION 52

- (Exam Topic 1)

A company stores sales transaction data in Amazon DynamoDB tables. To detect anomalous behaviors and respond quickly, all changes to the items stored in the DynamoDB tables must be logged within 30 minutes.

Which solution meets the requirements?

- A. Copy the DynamoDB tables into Apache Hive tables on Amazon EMR every hour and analyze them (or anomalous behavior)
- B. Send Amazon SNS notifications when anomalous behaviors are detected.
- C. Use AWS CloudTrail to capture all the APIs that change the DynamoDB table
- D. Send SNS notifications when anomalous behaviors are detected using CloudTrail event filtering.
- E. Use Amazon DynamoDB Streams to capture and send updates to AWS Lambda
- F. Create a Lambda function to output records to Amazon Kinesis Data Stream
- G. Analyze any anomalies with Amazon Kinesis Data Analytics
- H. Send SNS notifications when anomalous behaviors are detected.
- I. Use event patterns in Amazon CloudWatch Events to capture DynamoDB API call events with an AWS Lambda function as a target to analyze behavior
- J. Send SNS notifications when anomalous behaviors are detected.

Answer: C

Explanation:

[https://aws.amazon.com/blogs/database/dynamodb-streams-use-cases-and-design-patterns/#:~:text=DynamoDB DynamoDb Stream to capture DynamoDB update. And Kinesis Data Analytics for anomaly detection \(it uses AWS proprietary Random Cut Forest Algorithm\)](https://aws.amazon.com/blogs/database/dynamodb-streams-use-cases-and-design-patterns/#:~:text=DynamoDB DynamoDb Stream to capture DynamoDB update. And Kinesis Data Analytics for anomaly detection (it uses AWS proprietary Random Cut Forest Algorithm))

NEW QUESTION 57

- (Exam Topic 1)

A solutions architect is responsible (or redesigning a legacy Java application to improve its availability, data durability, and scalability. Currently, the application

runs on a single high-memory Amazon EC2 instance. It accepts HTTP requests from upstream clients, adds them to an in-memory queue, and responds with a 200 status. A separate application thread reads items from the queue, processes them, and persists the results to an Amazon RDS MySQL instance. The processing time for each item takes 90 seconds on average, most of which is spent waiting on external service calls, but the application is written to process multiple items in parallel.

Traffic to this service is unpredictable. During periods of high load, items may sit in the internal queue for over an hour while the application processes the backlog. In addition, the current system has issues with availability and data loss if the single application node fails.

Clients that access this service cannot be modified. They expect to receive a response to each HTTP request they send within 10 seconds before they will time out and retry the request.

Which approach would improve the availability and durability of the system while decreasing the processing latency and minimizing costs?

- A. Create an Amazon API Gateway REST API that uses Lambda proxy integration to pass requests to an AWS Lambda function
- B. Migrate the core processing code to a Lambda function and write a wrapper class that provides a handler method that converts the proxy events to the internal application data model and invokes the processing module.
- C. Create an Amazon API Gateway REST API that uses a service proxy to put items in an Amazon SQS queue
- D. Extract the core processing code from the existing application and update it to pull items from Amazon SQS instead of an in-memory queue
- E. Deploy the new processing application to smaller EC2 instances within an Auto Scaling group that scales dynamically based on the approximate number of messages in the Amazon SQS queue.
- F. Modify the application to use Amazon DynamoDB instead of Amazon RDS
- G. Configure Auto Scaling for the DynamoDB table
- H. Deploy the application within an Auto Scaling group with a scaling policy based on CPU utilization
- I. Back the in-memory queue with a memory-mapped file to an instance store volume and periodically write that file to Amazon S3.
- J. Update the application to use a Redis task queue instead of the in-memory queue
- K. Build a Docker container image for the application
- L. Create an Amazon ECS task definition that includes the application container and a separate container to host Redis
- M. Deploy the new task definition as an ECS service using AWS Fargate, and enable Auto Scaling.

Answer: B

Explanation:

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

NEW QUESTION 62

- (Exam Topic 1)

A team collects and routes behavioral data for an entire company. The company runs a Multi-AZ VPC environment with public subnets, private subnets, and an internet gateway. Each public subnet also contains a NAT gateway. Most of the company's applications read from and write to Amazon Kinesis Data Streams. Most of the workloads run in private subnets.

A solutions architect must review the infrastructure. The solutions architect needs to reduce costs and maintain the function of the applications. The solutions architect uses Cost Explorer and notices that the cost in the EC2-Other category is consistently high. A further review shows that NatGateway-Bytes charges are increasing the cost in the EC2-Other category.

What should the solutions architect do to meet these requirements?

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

Answer: D

Explanation:

<https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-endpoints-access.html> <https://aws.amazon.com/premiumsupport/knowledge-center/vpc-reduce-nat-gateway-transfer-costs/>

VPC endpoint policies enable you to control access by either attaching a policy to a VPC endpoint or by using additional fields in a policy that is attached to an IAM user, group, or role to restrict access to only occur via the specified VPC endpoint

NEW QUESTION 65

- (Exam Topic 1)

A solutions architect is evaluating the reliability of a recently migrated application running on AWS. The front end is hosted on Amazon S3 and accelerated by Amazon CloudFront. The application layer is running in a stateless Docker container on an Amazon EC2 On-Demand Instance with an Elastic IP address. The storage layer is a MongoDB database running on an EC2 Reserved Instance in the same Availability Zone as the application layer.

Which combination of steps should the solutions architect take to eliminate single points of failure with minimal application code changes? (Select TWO.)

- A. Create a REST API in Amazon API Gateway and use AWS Lambda functions as the application layer.
- B. Create an Application Load Balancer and migrate the Docker container to AWS Fargate.
- C. Migrate the storage layer to Amazon DynamoDB.
- D. Migrate the storage layer to Amazon DocumentDB (with MongoDB compatibility).
- E. Create an Application Load Balancer and move the storage layer to an EC2 Auto Scaling group.

Answer: BD

Explanation:

https://aws.amazon.com/documentdb/?nc1=h_ls

<https://aws.amazon.com/blogs/containers/using-alb-ingress-controller-with-amazon-eks-on-fargate/>

NEW QUESTION 69

- (Exam Topic 1)

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

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

Answer: D

NEW QUESTION 73

- (Exam Topic 1)

A company has a project that is launching Amazon EC2 instances that are larger than required. The project's account cannot be part of the company's organization in AWS Organizations due to policy restrictions to keep this activity outside of corporate IT. The company wants to allow only the launch of t3.small EC2 instances by developers in the project's account. These EC2 instances must be restricted to the us-east-2 Region. What should a solutions architect do to meet these requirements?

- A. Create a new developer account
- B. Move all EC2 instances, users, and assets into us-east-2. Add the account to the company's organization in AWS Organization
- C. Enforce a tagging policy that denotes Region affinity.
- D. Create an SCP that denies the launch of all EC2 instances except t3.small EC2 instances in us-east-2. Attach the SCP to the project's account.
- E. Create and purchase a t3.small EC2 Reserved Instance for each developer in us-east-2. Assign each developer a specific EC2 instance with their name as the tag.
- F. Create an IAM policy that allows the launch of only t3.small EC2 instances in us-east-2. Attach the policy to the roles and groups that the developers use in the project's account.

Answer: D

NEW QUESTION 78

- (Exam Topic 1)

A company hosts a web application that runs on a group of Amazon EC2 instances that are behind an Application Load Balancer (ALB) in a VPC. The company wants to analyze the network payloads to reverse-engineer a sophisticated attack of the application. Which approach should the company take to achieve this goal?

- A. Enable VPC Flow Log
- B. Store the flow logs in an Amazon S3 bucket for analysis.
- C. Enable Traffic Mirroring on the network interface of the EC2 instance
- D. Send the mirrored traffic to a target for storage and analysis.
- E. Create an AWS WAF web ACL
- F. and associate it with the ALB
- G. Configure AWS WAF logging.
- H. Enable logging for the ALB
- I. Store the logs in an Amazon S3 bucket for analysis.

Answer: A

NEW QUESTION 79

- (Exam Topic 1)

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

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

Answer: B

Explanation:

When EC2 instances reach third-party API through internet, their private 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-amaz>

NEW QUESTION 84

- (Exam Topic 1)

A company requires that all internal application connectivity use private IP addresses. To facilitate this policy, a solutions architect has created interface endpoints to connect to AWS public services. Upon testing, the solutions architect notices that the service names are resolving to public IP addresses, and that internal services cannot connect to the interface endpoints.

Which step should the solutions architect take to resolve this issue?

- A. Update the subnet route table with a route to the interface endpoint.
- B. Enable the private DNS option on the VPC attributes.
- C. Configure the security group on the interface endpoint to allow connectivity to the AWS services.
- D. Configure an Amazon Route 53 private hosted zone with a conditional forwarder for the internal application.

Answer: C

Explanation:

<https://docs.aws.amazon.com/vpc/latest/privatelink/vpce-interface.html>

NEW QUESTION 89

- (Exam Topic 1)

A company runs a popular public-facing ecommerce website. Its user base is growing quickly from a local market to a national market. The website is hosted in an on-premises data center with web servers and a MySQL database. The company wants to migrate its workload to AWS. A solutions architect needs to create a solution to:

- Improve security
- Improve reliability
- Improve availability
- Reduce latency
- Reduce maintenance

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

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

Answer: ABE

NEW QUESTION 91

- (Exam Topic 1)

A company wants to control its cost of Amazon Athena usage. The company has allocated a specific monthly budget for Athena usage. A solutions architect must design a solution that will prevent the company from exceeding the budgeted amount.

Which solution will meet these requirements?

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

Answer: D

NEW QUESTION 92

- (Exam Topic 1)

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

Which design should the solutions architect use?

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

- M. Use AWS CloudFormation StackSets to create the stacks in both Regions using Auto Scaling groups for the web and application tier
- N. Asynchronously replicate static content between Regions using Amazon S3 cross-Region replicatio
- O. Use Amazon CloudFront with static files in Amazon S3, and multi-Region origins for the front-end web tie
- P. Use Amazon DynamoD8 tables in each Region with scheduled backups to Amazon S3.

Answer: A

NEW QUESTION 96

- (Exam Topic 1)

A team collects and routes behavioral data for an entire company The company runs a Multi-AZ VPC environment with public subnets, private subnets, and in internet gateway Each public subnet also contains a NAT gateway Most of the company's applications read from and write to Amazon Kinesis Data Streams. Most of the workloads 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 Log
- B. Use Amazon Athena to analyze the logs for traffic that can be remove
- C. Ensure that security groups are Mocking traffic that is responsible for high costs.
- D. Add an interface VPC endpoint for Kinesis Data Streams to the VP
- E. Ensure that applications have the correct IAM permissions to use the interface VPC endpoint.
- F. 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
- G. Add an interface VPC endpoint for Kinesis Data Streams to the VP
- H. Ensure that the VPC endpoint policy allows traffic from the applications.

Answer: D

Explanation:

<https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-endpoints-access.html>

<https://aws.amazon.com/premiumsupport/knowledge-center/vpc-reduce-nat-gateway-transfer-costs/>

VPC endpoint policies enable you to control access by either attaching a policy to a VPC endpoint or by using additional fields in a policy that is attached to an IAM user, group, or role to restrict access to only occur via the specified VPC endpoint

NEW QUESTION 98

- (Exam Topic 1)

A company has an internal application running on AWS that is used to track and process shipments in the company's warehouse. Currently, after the system receives an order, it emails the staff the information needed to ship a package. Once the package is shipped, the staff replies to the email and the order is marked as shipped.

The company wants to stop using email in the application and move to a serverless application model. Which architecture solution meets these requirements?

- A. Use AWS Batch to configure the different tasks required lo ship a packag
- B. Have AWS Batch trigger an AWS Lambda function that creates and prints a shipping labe
- C. Once that label is scanne
- D. as it leaves the warehouse, have another Lambda function move the process to the next step in the AWS Batch job.B.
- E. When a new order is created, store the order information in Amazon SQ
- F. Have AWS Lambda check the queue every 5 minutes and process any needed wor
- G. When an order needs to be shipped, have Lambda print the label in the warehous
- H. Once the label has been scanned, as it leaves the warehouse, have an Amazon EC2 instance update Amazon SOS.
- I. Update the application to store new order information in Amazon DynamoD
- J. When a new order is created, trigger an AWS Step Functions workflow, mark the orders as "in progress," and print a package label to the warehous
- K. Once the label has been scanned and fulfilled, the application will trigger an AWS Lambda function that will mark the order as shipped and complete the workflow.
- L. Store new order information in Amazon EF
- M. Have instances pull the new information from the NFS and send that information to printers in the warehous
- N. Once the label has been scanned, as it leaves the warehouse, have Amazon API Gateway call the instances to remove the order information from Amazon EFS.

Answer: C

NEW QUESTION 101

- (Exam Topic 1)

A company wants to migrate a 30 TB Oracle data warehouse from on premises to Amazon Redshift The company used the AWS Schema Conversion Tool (AWS SCT) to convert the schema of the existing data warehouse to an Amazon Redshift schema The company also used a migration assessment report to identify manual tasks to complete.

The company needs to migrate the data to the new Amazon Redshift cluster during an upcoming data freeze period of 2 weeks The only network connection between the on-premises data warehouse and AWS is a 50 Mops internet connection

Which migration strategy meets these requirements?

- A. Create an AWS Database Migration Service (AWS DMS) replication instanc
- B. Authorize the public IP address of the replication instance to reach the data warehouse through the corporate firewall Create a migration task to run at the beginning of the data freeze period.
- C. Install the AWS SCT extraction agents on the on-premises server
- D. Define the extract, upload, and copy tasks to send the data to an Amazon S3 bucke
- E. Copy the data into the Amazon Redshift cluste
- F. Run the tasks at the beginning of the data freeze period.
- G. install the AWS SCT extraction agents on the on-premises server
- H. Create a Site-to-Site VPN connection Create an AWS Database Migration Service (AWS DMS) replication instance that is the appropriate size Authorize the IP address of the replication instance to be able to access the on-premises data warehouse through the VPN connection

I. Create a job in AWS Snowball Edge to import data into Amazon S3 Install AWS SCT extraction agents on the on-premises servers Define the local and AWS Database Migration Service (AWS DMS) tasks to send the data to the Snowball Edge device When the Snowball Edge device is returned to AWS and the data is available in Amazon S3, run the AWS DMS subtask to copy the data to Amazon Redshift.

Answer: D

Explanation:

AWS Database Migration Service (AWS DMS) can use Snowball Edge and Amazon S3 to migrate large databases more quickly than by other methods

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_LargeDBs.html

https://www.calctool.org/CALC/prof/computing/transfer_time

NEW QUESTION 106

- (Exam Topic 1)

A company runs an application on AWS. An AWS Lambda function uses credentials to authenticate to an Amazon RDS for MySQL DB instance. A security risk assessment identified that these credentials are not frequently rotated. Also, encryption at rest is not enabled for the DB instance. The security team requires that both of these issues be resolved.

Which strategy should a solutions architect recommend to remediate these security risks?

- A. Configure the Lambda function to store and retrieve the database credentials in AWS Secrets Manager and enable rotation of the credential
- B. Take a snapshot of the DB instance and encrypt a copy of that snapshot
- C. Replace the DB instance with a new DB instance that is based on the encrypted snapshot.
- D. Enable IAM DB authentication on the DB instance
- E. Grant the Lambda execution role access to the DB instance
- F. Modify the DB instance and enable encryption.
- G. Enable IAM DB authentication on the DB instance
- H. Grant the Lambda execution role access to the DB instance
- I. Create an encrypted read replica of the DB instance
- J. Promote the encrypted read replica to be the new primary node.
- K. Configure the Lambda function to store and retrieve the database credentials as encrypted AWS Systems Manager Parameter Store parameter
- L. Create another Lambda function to automatically rotate the credential
- M. Create an encrypted read replica of the DB instance
- N. Promote the encrypted read replica to be the new primary node.

Answer: A

Explanation:

Parameter store can store DB credentials as secure string but CANNOT rotate secrets, hence, go with A + Cannot enable encryption on existing MySQL RDS instance, must create a new encrypted one from unencrypted snapshot.

<https://aws.amazon.com/blogs/security/rotate-amazon-rds-database-credentials-automatically-with-aws-secrets-> Encrypting a unencrypted instance of DB or creating a encrypted replica of an unencrypted DB instance are not possible Hence A is the only solution possible.

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html#Overview.Encryption>.

NEW QUESTION 109

- (Exam Topic 1)

A company needs to implement a patching process for its servers. The on-premises servers and Amazon EC2 instances use a variety of tools to perform patching. Management requires a single report showing the patch status of all the servers and instances.

Which set of actions should a solutions architect take to meet these requirements?

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

Answer: A

Explanation:

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

NEW QUESTION 110

- (Exam Topic 2)

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


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

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

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

Answer: B

NEW QUESTION 111

- (Exam Topic 2)

A company is migrating its marketing website and content management system from an on-premises data center to AWS. The company wants the AWS application to be deployed in a VPC with Amazon EC2 instances used for the web servers and an Amazon RDS instance for the database.

The company has a runbook document that describes the installation process of the on-premises system. The company would like to base the AWS system on the processes referenced in the runbook document. The runbook document describes the installation and configuration of the operating systems, network settings, the website, and content management system software on the servers. After the migration is complete, the company wants to be able to make changes quickly to take advantage of other AWS features.

How can the application and environment be deployed and automated in AWS, while allowing for future changes?

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

Answer: D

NEW QUESTION 115

- (Exam Topic 2)

A company that develops consumer electronics with offices in Europe and Asia has 60 TB of software images stored on premises in Europe. The company wants to transfer the images to an Amazon S3 bucket in the ap-northeast-1 Region. New software images are created daily and must be encrypted in transit. The company needs a solution that does not require custom development to automatically transfer all existing and new software images to Amazon S3.

What is the next step in the transfer process?

- A. Deploy an AWS DataSync agent and configure a task to transfer the images to the S3 bucket.
- B. Configure Amazon Kinesis Data Firehose to transfer the images using S3 Transfer Acceleration.
- C. Use an AWS Snowball device to transfer the images with the S3 bucket as the target.
- D. Transfer the images over a Site-to-Site VPN connection using the S3 API with multipart upload.

Answer: A

NEW QUESTION 119

- (Exam Topic 2)

A company has developed a new billing application that will be released in two weeks. Developers are testing the application running on 10 EC2 instances managed by an Auto Scaling group in subnet 172.31.0.0/24 within VPC A with CIDR block 172.31.0.0/16. The developers noticed connection timeout errors in the application logs while connecting to an Oracle database running on an Amazon EC2 instance in the same region within VPC B with CIDR block 172.50.0.0/16. The IP of the database instance is hard-coded in the application instances.

Which recommendations should a solutions architect present to the developers to solve the problem in a secure way with minimal maintenance and overhead?

- A. Disable the SrcDestCheck attribute for all instances running the application and Oracle Database. Change the default route of VPC A to point ENI of the Oracle Database that has an IP address assigned within the range of 172.50.0.0/16
- B. Create and attach internet gateways for both VPC
- C. Configure default routes to the internet gateways for both VPC
- D. Assign an Elastic IP for each Amazon EC2 instance in VPC A
- E. Create a VPC peering connection between the two VPCs and add a route to the routing table of VPC A that points to the IP address range of 172.50.0.0/16
- F. Create an additional Amazon EC2 instance for each VPC as a customer gateway; create one virtual private gateway (VGW) for each VP
- G. configure an end-to-end VPC, and advertise the routes for 172.50.0.0/16

Answer: C

NEW QUESTION 124

- (Exam Topic 2)

A company is running a two-tier web-based application in an on-premises data center. The application layer consists of a single server running a stateful application. The application connects to a PostgreSQL database running on a separate server. The application's user base is expected to grow significantly, so the company is migrating the application and database to AWS. The solution will use Amazon Aurora PostgreSQL, Amazon EC2 Auto Scaling, and Elastic Load Balancing.

Which solution will provide a consistent user experience that will allow the application and database tiers to scale?

- A. Enable Aurora Auto Scaling for Aurora Replica
- B. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled
- C. Enable Aurora Auto Scaling for Aurora writer
- D. Use an Application Load Balancer with the round robin routing algorithm and sticky sessions enabled
- E. Aurora Auto Scaling for Aurora Replica
- F. Use an Application Load Balancer with the round robin routing algorithm and sticky sessions enabled.
- G. Aurora Auto Scaling for Aurora writer
- H. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.

Answer: C

NEW QUESTION 127

- (Exam Topic 2)

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

Which strategy should the solutions architect use?

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

Answer: C

NEW QUESTION 132

- (Exam Topic 2)

A company wants to migrate its website from an on-premises data center onto AWS. At the same time, it wants to migrate the website to a containerized microservice-based architecture to improve the availability and cost efficiency. The company's security policy states that privileges and network permissions must be configured according to best practice, using least privilege.

A solutions architect must create a containerized architecture that meets the security requirements and has deployed the application to an Amazon ECS cluster. What steps are required after the deployment to meet the requirements? (Select TWO.)

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

Answer: BE

NEW QUESTION 133

- (Exam Topic 2)

A company uses AWS Organizations with a single OU named Production to manage multiple accounts. All accounts are members of the Production OU. Administrators use deny list SCPs in the root of the organization to manage access to restricted services.

The company recently acquired a new business unit and invited the new unit's existing AWS account to the

organization. Once onboarded, the administrators of the new business unit discovered that they are not able to update existing AWS Config rules to meet the company's policies.

Which option will allow administrators to make changes and continue to enforce the current policies without introducing additional long-term maintenance?

- A. Remove the organization's root SCPs that limit access to AWS Config. Create AWS Service Catalog products for the company's standard AWS Config rules and deploy them throughout the organization, including the new account.
- B. Create a temporary OU named Onboarding for the new account. Apply an SCP to the Onboarding OU to allow AWS Config actions. Move the new account to the Production OU when adjustments to AWS Config are complete.
- C. Convert the organization's root SCPs from deny list SCPs to allow list SCPs to allow the required services only. Temporarily apply an SCP to the organization's root that allows AWS Config actions for principals only in the new account.
- D. Create a temporary OU named Onboarding for the new account. Apply an SCP to the Onboarding OU to allow AWS Config action.
- E. Move the organization's root SCP to the Production OU.
- F. Move the new account to the Production OU when adjustments to AWS Config are complete.

Answer: D

NEW QUESTION 136

- (Exam Topic 2)

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

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

Answer: C

NEW QUESTION 139

- (Exam Topic 2)

A retail company is running an application that stores invoice files in an Amazon S3 bucket and metadata about the files in an Amazon DynamoDB table. The application software runs in both us-east-1 and eu-west-1. The S3 bucket and DynamoDB table are in us-east-1. The company wants to protect itself from data corruption and loss of connectivity to either Region. Which option meets these requirements?

- A. Create a DynamoDB global table to replicate data between us-east-1 and eu-west-1. Enable continuous backup on the DynamoDB table in us-east-1. Enable versioning on the S3 bucket.
- B. Create an AWS Lambda function triggered by Amazon CloudWatch Events to make regular backups of the DynamoDB table. Set up S3 cross-region replication from us-east-1 to eu-west-1. Set up MFA delete on the S3 bucket in us-east-1.
- C. Create a DynamoDB global table to replicate data between us-east-1 and eu-west-1. Enable versioning on the S3 bucket. Implement strict ACLs on the S3 bucket.
- D. Create a DynamoDB global table to replicate data between us-east-1 and eu-west-1. Enable continuous backup on the DynamoDB table in us-east-1. Set up S3 cross-region replication from us-east-1 to eu-west-1.

Answer: B

NEW QUESTION 141

- (Exam Topic 2)

A Solutions Architect is constructing a containerized .NET Core application for AWS Fargate. The application's backend needs a high-availability version of Microsoft SQL Server. All application levels must be extremely accessible. The credentials associated with the SQL Server connection string should not be saved to disk inside the .NET Core front-end containers. Which tactics should the Solutions Architect use to achieve these objectives?

- A. Set up SQL Server to run in Fargate with Service Auto Scaling.
- B. Create an Amazon ECS task execution role that allows the Fargate task definition to get the secret value for the credentials to SQL Server running in Fargate.
- C. Specify the ARN of the secret in AWS Secrets Manager in the secrets section of the Fargate task definition so the sensitive data can be injected into the containers as environment variables on startup for reading into the application to construct the connection string.
- D. Set up the .NET Core service using Service Auto Scaling behind an Application Load Balancer in multiple Availability Zones.
- E. Create a Multi-AZ deployment of SQL Server on Amazon RDS.
- F. Create a secret in AWS Secrets Manager for the credentials to the RDS database.
- G. Create an Amazon ECS task execution role that allows the Fargate task definition to get the secret value for the credentials to the RDS database in Secrets Manager.
- H. Specify the ARN of the secret in Secrets Manager in the secrets section of the Fargate task definition so the sensitive data can be injected into the containers as environment variables on startup for reading into the application to construct the connection string.
- I. Set up the .NET Core service in Fargate using Service Auto Scaling behind an Application Load Balancer in multiple Availability Zones.
- J. Create an Auto Scaling group to run SQL Server on Amazon EC2. Create a secret in AWS Secrets Manager for the credentials to SQL Server running on EC2. Specify the ARN of the secret in Secrets Manager in the secrets section of the Fargate task definition so the sensitive data can be injected into the containers as environment variables on startup for reading into the application to construct the connection string.
- K. Set up the .NET Core service using Service Auto Scaling behind an Application Load Balancer in multiple Availability Zones.
- L. Create a Multi-AZ deployment of SQL Server on Amazon RDS.
- M. Create a secret in AWS Secrets Manager for the credentials to the RDS database.
- N. Create non-persistent empty storage for the .NET Core containers in the Fargate task definition to store the sensitive information.
- O. Create an Amazon ECS task execution role that allows the Fargate task definition to get the secret value for the credentials to the RDS database in Secrets Manager.
- P. Specify the ARN of the secret in Secrets Manager in the secrets section of the Fargate task definition so the sensitive data can be written to the non-persistent empty storage on startup for reading into the application to construct the connection string.
- Q. Set up the .NET Core service using Service Auto Scaling behind an Application Load Balancer in multiple Availability Zones.

Answer: B

Explanation:

Secrets Manager natively supports SQL Server on RDS. No real need to create additional 'ephemeral storage' to fetch credentials, as these can be injected to containers as environment variables. <https://aws.amazon.com/premiumsupport/knowledge-center/ecs-data-security-container-task/>

NEW QUESTION 143

- (Exam Topic 2)

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 tool to create dashboards to observe events in near-real time.
- Support for semi-structured JSON data and dynamic schemas.

Which combination of components will enable the company to create a monitoring solution that will satisfy these requirements? (Select TWO.)

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

Answer: DE

NEW QUESTION 146

- (Exam Topic 2)

A solutions architect has been assigned to migrate a 50 TB Oracle data warehouse that contains sales data from on-premises to Amazon Redshift. Major updates to the sales data occur on the final calendar day of the month. For the remainder of the month, the data warehouse only receives minor daily updates and is primarily used for reading and reporting. Because of this, the migration process must start on the first day of the month and must be complete before the next set of updates occur. This provides approximately 30 days to complete the migration and ensure that the minor daily changes have been synchronized with the Amazon Redshift data warehouse. Because the migration cannot impact normal business network operations, the bandwidth allocated to the migration for moving data over the internet is 50 Mbps. The company wants to keep data migration costs low.

Which steps will allow the solutions architect to perform the migration within the specified timeline?

- A. Install Oracle database software on an Amazon EC2 instance. Configure VPN connectivity between AWS and the company's data center. Configure the Oracle database running on Amazon EC2 to join the Oracle Real Application Clusters (RAC). When the Oracle database on Amazon EC2 finishes synchronizing, create an AWS DMS ongoing replication task to migrate the data from the Oracle database on Amazon EC2 to Amazon Redshift. Verify the data migration is complete and perform the cut over to Amazon Redshift.
- B. Create an AWS Snowball import job. Export a backup of the Oracle data warehouse. Copy the exported data to the Snowball device. Return the Snowball device to AWS. Create an Amazon RDS for Oracle database and restore the backup file to that RDS instance. Create an AWS DMS task to migrate the data from the RDS for Oracle database to Amazon Redshift. Copy daily incremental backups from Oracle in the data center to the RDS for Oracle database over the internet. Verify the data migration is complete and perform the cut over to Amazon Redshift.
- C. Install Oracle database software on an Amazon EC2 instance. To minimize the migration time, configure VPN connectivity between AWS and the company's data center by provisioning a 1 Gbps AWS Direct Connect connection. Configure the Oracle database running on Amazon EC2 to be a read replica of the data center Oracle database. Start the synchronization process between the company's on-premises data center and the Oracle database on Amazon EC2. When the Oracle database on Amazon EC2 is synchronized with the on-premises database, create an AWS DMS ongoing replication task from the Oracle database read replica that is running on Amazon EC2 to Amazon Redshift. Verify the data migration is complete and perform the cut over to Amazon Redshift.
- D. Create an AWS Snowball import job.
- E. Configure a server in the company's data center with an extraction agent.
- F. Use AWS SCT to manage the extraction agent and convert the Oracle schema to an Amazon Redshift schema.
- G. Create a new project in AWS SCT using the registered data extraction agent.
- H. Create a local task and an AWS DMS task in AWS SCT with replication of ongoing changes.
- I. Copy data to the Snowball device and return the Snowball device to AWS.
- J. Allow AWS DMS to copy data from Amazon S3 to Amazon Redshift.
- K. Verify that the data migration is complete and perform the cut over to Amazon Redshift.

Answer: D

Explanation:

Create an AWS Snowball import job. Configure a server in the company's data center with an extraction agent. Use AWS SCT to manage the extraction agent and convert the Oracle schema to an Amazon Redshift schema. Create a new project in AWS SCT using the registered data extraction agent. Create a local task and an AWS DMS task in AWS SCT with replication of ongoing changes. Copy data to the Snowball device and return the Snowball device to AWS. Allow AWS DMS to copy data from Amazon S3 to Amazon Redshift. Verify that the data migration is complete and perform the cut over to Amazon Redshift.

<https://aws.amazon.com/getting-started/hands-on/migrate-oracle-to-amazon-redshift/>

NEW QUESTION 148

- (Exam Topic 2)

A software development company has multiple engineers who are working remotely. The company is running Active Directory Domain Services (AD DS) on an Amazon EC2 instance. The company's security policy states that all internal, nonpublic services that are deployed in a VPC must be accessible through a VPN. Multi-factor authentication (MFA) must be used for access to a VPN.

What should a solution architect do to meet these requirements?

- A. Create an AWS Site-to-Site VPN connection. Configure integration between a VPN and AD DS.
- B. Use an Amazon WorkSpaces client with MFA support enabled to establish a VPN connection.
- C. Create an AWS Client VPN endpoint. Create an AD Connector directory for integration with AD DS. Enable MFA for AD Connector. Use AWS Client VPN to establish a VPN connection.
- D. Create multiple AWS Site-to-Site VPN connections by using AWS VPN CloudHub. Configure integration between AWS VPN CloudHub and AD DS. Use AWS CloudHub to establish a VPN connection.
- E. Create an Amazon WorkLink endpoint. Configure integration between Amazon WorkLink and AD DS.
- F. Enable MFA in Amazon WorkLink. Use AWS Client VPN to establish a VPN connection.

Answer: B

NEW QUESTION 153

- (Exam Topic 2)

A company is migrating its infrastructure to the AWS Cloud. The company must comply with a variety of regulatory standards for different projects. The company needs a multi-account environment.

A solutions architect needs to prepare the baseline infrastructure. The solution must provide a consistent baseline of management and security but it must allow flexibility for different compliance requirements within various AWS accounts. The solution also needs to integrate with the existing on-premises Active Directory Federation Services (AD FS) server.

Which solution meets these requirements with the LEAST amount of operational overhead?

- A. Create an organization in AWS Organizations. Create a single SCP for least privilege access across all accounts. Create a single OU for all accounts. Configure an IAM identity provider for federation with the on-premises AD FS server. Configure a central logging account with a defined process for log generating services to send log events to the central account.
- B. Enable AWS Config in the central account with conformance packs for all accounts.
- C. Create an organization in AWS Organizations. Enable AWS Control Tower on the organization.
- D. Review included guardrails for SCP.
- E. Check AWS Config for areas that require additions. Add OUs as necessary. Connect AWS Single Sign-On to the on-premises AD FS server.
- F. Create an organization in AWS Organizations. Create SCPs for least privilege access. Create an OU structure, and use it to group AWS accounts. Connect AWS Single Sign-On to the on-premises AD FS server.
- G. Configure a central logging account with a defined process for log generating services to send log events to the central account. Enable AWS Config in the central account with aggregators and conformance packs.
- H. Create an organization in AWS Organizations. Enable AWS Control Tower on the organization. Review included guardrails for SCP.
- I. Check AWS Config for areas that require additions. Configure an IAM identity provider for federation with the on-premises AD FS server.

Answer: A

NEW QUESTION 158

- (Exam Topic 2)

A company has multiple business units. Each business unit has its own AWS account and runs a single website within that account. The company also has a single logging account. Logs from each business unit website are aggregated into a single Amazon S3 bucket in the logging account. The S3 bucket policy provides each business unit with access to write data into the bucket and requires data to be encrypted.

The company needs to encrypt logs uploaded into the bucket using a Single AWS Key Management Service (AWS KMS) CMK. The CMK that protects the data must be rotated once every 365 days.

Which strategy is the MOST operationally efficient for the company to use to meet these requirements?

- A. Create a customer managed CMK in the logging account. Update the CMK key policy to provide access to the logging account only. Manually rotate the CMK every 365 days.
- B. Create a customer managed CMK in the logging account.
- C. Update the CMK key policy to provide access to the logging account and business unit account.
- D. Enable automatic rotation of the CMK.
- E. Use an AWS managed CMK in the logging account.
- F. Update the CMK key policy to provide access to the logging account and business unit accounts. Manually rotate the CMK every 365 days.
- G. Use an AWS managed CMK in the logging account. Update the CMK key policy to provide access to the logging account only.
- H. Enable automatic rotation of the CMK.

Answer: A

NEW QUESTION 160

- (Exam Topic 2)

A large company recently experienced an unexpected increase in Amazon RDS and Amazon DynamoDB costs. The company needs to increase visibility into details of AWS Billing and Cost Management. There are various accounts associated with AWS Organizations, including many development and production accounts. There is no consistent tagging strategy across the organization, but there are guidelines in place that require all infrastructure to be deployed using AWS CloudFormation with consistent tagging. Management requires cost center numbers and project ID numbers for all existing and future DynamoDB tables and RDS instances.

Which strategy should the solutions architect provide to meet these requirements?

- A. Use Tag Editor to tag existing resources. Create cost allocation tags to define the cost center and project ID and allow 24 hours for tags to propagate to existing resources.
- B. Use an AWS Config rule to alert the finance team of untagged resources. Create a centralized AWS Lambda based solution to tag untagged RDS databases and DynamoDB resources every hour using a cross-account role.
- C. Use Tag Editor to tag existing resources. Create cost allocation tags to define the cost center and project ID. Use SCPs to restrict resource creation that do not have the cost center and project ID on the resource.
- D. Create cost allocation tags to define the cost center and project ID and allow 24 hours for tags to propagate to existing resources. Update existing federated roles to restrict privileges to provision resources that do not include the cost center and project ID on the resource.

Answer: B

NEW QUESTION 165

- (Exam Topic 2)

A company operates quick-service restaurants. The restaurants follow a predictable model with high sales traffic for 4 hours daily. Sales traffic is lower outside of those peak hours.

The point of sale and management platform is deployed in the AWS Cloud and has a backend that is based on Amazon DynamoDB. The database table uses provisioned throughput mode with 100,000 RCUs and 80,000 WCUs to match known peak resource consumption.

The company wants to reduce its DynamoDB cost and minimize the operational overhead for the IT staff. Which solution meets these requirements MOST cost-effectively?

- A. Reduce the provisioned RCUs and WCUs.
- B. Change the DynamoDB table to use on-demand capacity.
- C. Enable DynamoDB auto scaling for the table.
- D. Purchase 1-year reserved capacity that is sufficient to cover the peak load for 4 hours each day.

Answer: C

NEW QUESTION 166

- (Exam Topic 2)

A company owns a chain of travel agencies and is running an application in the AWS Cloud. Company employees use the application to search (or Information about travel destinations. Destination content is updated four times each year.

Two fixed Amazon EC2 instances serve the application. The company uses an Amazon Route 53 public hosted zone with a multivalue record of travel.example.com that returns the Elastic IP addresses for the EC2 instances. The application uses Amazon DynamoDB as its primary data store. The company uses a self-hosted Redis instance as a caching solution.

During content updates, the load on the EC2 instances and the caching solution increases drastically. This increased load has led to downtime on several occasions. A solutions architect must update the application so that the application is highly available and can handle the load that is generated by the content updates.

Which solution will meet these requirements?

- A. Set up DynamoDB Accelerator (DAX) as in-memory cach
- B. Update the application to use DA
- C. Create an Auto Scaling group for the EC2 instance
- D. Create an Application Load Balancer (ALB). Set the Auto Scaling group as a target for the AL
- E. Update the Route 53 record to use a simple routing policy that targets the ALB's DNS alia
- F. Configure scheduled scaling for the EC2 instances before the content updates.
- G. Set up Amazon ElastiCache for Redi
- H. Update the application to use ElastiCach
- I. Create an Auto Scaling group for the EC2 instance
- J. Create an AmazonCloudFront distnbutio
- K. and set the Auto Scaling group as an origin for the distributio
- L. Update the Route 53 record to use a simple routing policy that targets the CloudFront distribution's DNS alias Manually scale up EC2 instances before the content updates
- M. Set up Amazon ElastiCache for Memcache
- N. Update the application to use ElastiCach
- O. Create an Auto Scaling group for the EC2 instances Create an Application Load Balancer (ALB). Set the Auto Scaling group as a target for the AL
- P. Update the Route 53 record to use a simple routing policy that targets the ALB's DNS alia
- Q. Configure scheduled scaling for the application before the content updates.
- R. Set up DynamoDB Accelerator (DAX) as in-memory cach
- S. Update the application to use DA
- T. Create an Auto Scaling group for the EC2 instance
- . Create an Amazon CloudFront distribution, and set the Auto Scaling group as an origin for the distributio
- . Update the Route 53 record to use a simple routing policy that targets the CloudFront distribution's DNS alia
- . Manually scale up EC2 instances before the content updates.

Answer: B

NEW QUESTION 170

- (Exam Topic 2)

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

Answer: ABE

NEW QUESTION 175

- (Exam Topic 2)

A financial services company loaded millions of historical stock trades into an Amazon DynamoDB table. The table uses on-demand capacity mode. Once each day at midnight, a few million new records are loaded into the table. Application read activity against the table happens in bursts throughout the day, and a limited set of keys are repeatedly looked up. The company needs to reduce costs associated with DynamoDB.

Which strategy should a solutions architect recommend to meet this requirement?

- A. Deploy an Amazon ElastiCache cluster in front of the DynamoDB table.
- B. Deploy DynamoDB Accelerator (DAX). Configure DynamoDB auto scaling. Purchase Savings Plans in Cost Explorer.
- C. Use provisioned capacity mode. Purchase Savings Plans in Cost Explorer.
- D. Deploy DynamoDB Accelerator (DAX). Use provisioned capacity mode. Configure DynamoDB auto scaling.

Answer: D

NEW QUESTION 180

- (Exam Topic 2)

A company is running a critical application that uses an Amazon RDS for MySQL database to store data. The RDS DB instance is deployed in Multi-AZ mode. A recent RDS database failover test caused a 40-second outage to the application. A solutions architect needs to design a solution to reduce the outage time to less than 20 seconds.

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

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

Answer: ABF

NEW QUESTION 182

- (Exam Topic 2)

A company recently deployed a new application that runs on a group of Amazon EC2 Linux instances in a VPC. In a peered VPC, the company launched an EC2 Linux instance that serves as a bastion host. The security group of the application instances allows access only on TCP port 22 from the private IP of the bastion host. The security group of the bastion host allows access to TCP port 22 from 0.0.0.0/0 so that system administrators can use SSH to remotely log in to the application instances from several branch offices.

While looking through operating system logs on the bastion host, a cloud engineer notices thousands of failed SSH logins to the bastion host from locations around the world. The cloud engineer wants to change how remote access is granted to the application instances and wants to meet the following requirements:

- Eliminate brute-force SSH login attempts
- Retain a log of commands run during an SSH session
- Retain the ability to forward ports

Which solution meets these requirements for remote access to the application instances?

- A. Configure the application instances to communicate with AWS Systems Manager. Grant access to the system administrators to use Session Manager to establish a session with the application instances. Terminate the bastion host.
- B. Update the security group of the bastion host to allow traffic from only the public IP addresses of the branch offices.
- C. Configure an AWS Client VPN endpoint and provision each system administrator with a certificate to establish a VPN connection to the application VPC. Update the security group of the application instances to allow traffic from only the Client VPN IPv4 CIDR.
- D. Terminate the bastion host.
- E. Configure the application instances to communicate with AWS Systems Manager.
- F. Grant access to the system administrators to issue commands to the application instances by using Systems Manager Run Command.
- G. Terminate the bastion host.

Answer: A

Explanation:

"Session Manager removes the need to open inbound ports, manage SSH keys, or use bastion hosts" Ref: <https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager.html>

NEW QUESTION 187

- (Exam Topic 2)

A company is creating a sequel for a popular online game. A large number of users from all over the world will play the game within the first week after launch. Currently, the game consists of the following components deployed in a single AWS Region:

- Amazon S3 bucket that stores game assets
- Amazon DynamoDB table that stores player scores

A solutions architect needs to design a Region solution that will reduce latency, improve reliability, and require the least effort to implement.

What should the solutions architect do to meet these requirements?

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

Answer: B

NEW QUESTION 189

- (Exam Topic 2)

A company is running its solution on AWS in a manually created VPC. The company is using AWS CloudFormation to provision other parts of the infrastructure. According to a new requirement, the company must manage all infrastructure in an automatic way.

What should the company do to meet this new requirement with the LEAST effort?

- A. Create a new AWS Cloud Development Kit (AWS CDK) stack that strictly provisions the existing VPC resources and configuration.
- B. Use AWS CDK to import the VPC into the stack and to manage the VPC.
- C. Create a CloudFormation stack set that creates the VPC.
- D. Use the stack set to import the VPC into the stack.
- E. Create a new CloudFormation template that strictly provisions the existing VPC resources and configuration.
- F. From the CloudFormation console, create a new stack by importing the existing resources.
- G. Create a new CloudFormation template that creates the VPC.
- H. Use the AWS Serverless Application Model (AWS SAM) CLI to import the VPC.

Answer: D

NEW QUESTION 192

- (Exam Topic 2)

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

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

Answer: B

NEW QUESTION 196

- (Exam Topic 2)

A car rental company has built a serverless REST API to provide data to its mobile app. The app consists of an Amazon API Gateway API with a Regional endpoint, AWS Lambda functions, and an Amazon Aurora MySQL Serverless DB cluster. The company recently opened the API to mobile apps of partners. A significant increase in the number of requests resulted in causing sporadic database memory errors. Analysis of the API traffic indicates that clients are making multiple HTTP GET requests for the same queries in a short period of time. Traffic is concentrated during business hours, with spikes around holidays and other events. The company needs to improve its ability to support the additional usage while minimizing the increase in costs associated with the solution. Which strategy meets these requirements?

- A. Convert the API Gateway Regional endpoint to an edge-optimized endpoint. Enable caching in the production stage.
- B. Implement an Amazon ElastiCache for Redis cache to store the results of the database calls. Modify the Lambda functions to use the cache.
- C. Modify the Aurora Serverless DB cluster configuration to increase the maximum amount of available memory.
- D. Enable throttling in the API Gateway production stage. Set the rate and burst values to limit the incoming calls.

Answer: A

NEW QUESTION 197

- (Exam Topic 2)

A company plans to refactor a monolithic application into a modern application designed to be deployed on AWS. The CI/CD pipeline needs to be upgraded to support the modern design for the application with the following requirements:

- It should allow changes to be released several times every hour.
 - * It should be able to roll back the changes as quickly as possible.
- Which design will meet these requirements?

- A. Deploy a CI-CD pipeline that incorporates AMIs to contain the application and their configurations. Deploy the application by replacing Amazon EC2 instances.
- B. Specify AWS Elastic Beanstalk to serve as a secondary environment as the deployment target for the CI/CD pipeline of the application.
- C. To deploy, swap the staging and production environment URLs.
- D. Use AWS Systems Manager to re-provision the infrastructure for each deployment. Update the Amazon EC2 user data to pull the latest code artifact from Amazon S3 and use Amazon Route 53 weighted routing to point to the new environment.
- E. Roll out application updates as part of an Auto Scaling event using a prebuilt AMI.
- F. Use new versions of the AMIs to add instances, and phase out all instances that use the previous AMI version with the configured termination policy during a deployment event.

Answer: B

Explanation:

It is the fastest when it comes to rollback and deploying changes every hour.

NEW QUESTION 199

- (Exam Topic 2)

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

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

Answer: A

NEW QUESTION 203

- (Exam Topic 2)

A new startup is running a serverless application using AWS Lambda as the primary source of compute. New versions of the application must be made available to a subset of users before deploying changes to all users. Developers should also have the ability to stop the deployment and have access to an easy rollback mechanism. A solutions architect decides to use AWS CodeDeploy to deploy changes when a new version is available. Which CodeDeploy configuration should the solutions architect use?

- A. A blue/green deployment
- B. A linear deployment
- C. A canary deployment

D. An all-at-once deployment

Answer: C

NEW QUESTION 208

- (Exam Topic 2)

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

Which tools or services should the solutions architect use to plan the cloud migration? (Select THREE.)

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

Answer: ADF

NEW QUESTION 210

- (Exam Topic 2)

A company is using multiple AWS accounts The DNS records are stored in a private hosted zone for Amazon Route 53 in Account A The company's applications and databases are running in Account B.

A solutions architect will deploy a two-tier application In a new VPC To simplify the configuration, the db.example.com CNAME record set for the Amazon RDS endpoint was created in a private hosted zone for Amazon Route 53.

During deployment, the application failed to start. Troubleshooting revealed that db.example.com is not resolvable on the Amazon EC2 instance The solutions architect confirmed that the record set was created correctly in Route 53.

Which combination of steps should the solutions architect take to resolve this issue? (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/etc/resolv.conf file
- C. Create an authorization to associate the private hosted zone in Account A with the new VPC in Account B
- D. Create a private hosted zone for the example.com domain in Account B Configure Route 53 replication between AWS accounts
- E. Associate a new VPC in Account B with a hosted zone in Account A
- F. Delete the association authorization in Account A.

Answer: CE

NEW QUESTION 214

- (Exam Topic 2)

A news company wants to implement an AWS Lambda function that calls an external API to receive new press releases every 10 minutes. The API provider is planning to use an IP address allow list to protect the API. so the news company needs to provide any public IP addresses that access the API. The company's current architecture includes a VPC with an internet gateway and a NAT gateway. A solutions architect must implement a static IP address for the Lambda function.

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

- A. Use the Elastic IP address that is associated with the NAT gateway for the IP address allow list.
- B. Assign an Elastic IP address to the Lambda function
- C. Use the Lambda function's Elastic IP address for the IP address allow list.
- D. Configure the Lambda function to launch in the private subnet of the VPC.
- E. Configure the Lambda function to launch in the public subnet of the VPC.
- F. Create a transit gateway
- G. Attach the VPC and the Lambda function to the transit gateway.

Answer: AC

NEW QUESTION 219

- (Exam Topic 2)

A company has an organization in AWS Organizations. The organization consists of a large number of AWS accounts that belong to separate business units. The company requires all Amazon EC2 instances to be provisioned with custom, hardened AMIs. The company wants a solution that provides each AWS account access to the AMIs

Which solution will meet these requirements with the MOST operational efficiency?

- A. Create the AMIs with EC2 Image Builder Create an AWS CodePipeline pipeline to share the AMIs across all AWS accounts.
- B. Deploy Jenkins on an EC2 instance Create jobs to create and share the AMIs across all AWS accounts.
- C. Create and share the AMIs with EC2 Image Builder Use AWS Service Catalog to configure a product that provides access to the AMIs across all AWS accounts.
- D. Create the AMIs with EC2 Image Builder Create an AWS Lambda function to share the AMIs across all AWS accounts.

Answer: C

NEW QUESTION 220

- (Exam Topic 2)

A flood monitoring agency has deployed more than 10,000 water-level monitoring sensors. Sensors send continuous data updates, and each update is less than 1 MB in size. The agency has a fleet of on-premises application servers. These servers receive updates from the sensors, convert the raw data into a human readable format, and write the results to an on-premises relational database server Data analysts then use simple SQL queries to monitor the data.

The agency wants to increase overall application availability and reduce the effort that is required to perform maintenance tasks. These maintenance tasks, which include updates and patches to the application servers, cause downtime. While an application server is down, data is lost from sensors because the remaining servers cannot handle the entire workload.

The agency wants a solution that optimizes operational overhead and costs. A solutions architect recommends the use of AWS IoT Core to collect the sensor data. What else should the solutions architect recommend to meet these requirements?

- A. Send the sensor data to Amazon Kinesis Data Firehose
- B. Use an AWS Lambda function to read the Kinesis Data Firehose data, convert it to .csv format, and insert it into an Amazon Aurora MySQL DB Instance
- C. Instruct the data analysts to query the data directly from the DB Instance.
- D. Send the sensor data to Amazon Kinesis Data Firehose
- E. Use an AWS Lambda function to read the Kinesis Data Firehose data, convert it to Apache Parquet format, and save it to an Amazon S3 bucket
- F. Instruct the data analysts to query the data by using Amazon Athena.
- G. Send the sensor data to an Amazon Kinesis Data Analytics application to convert the data to csv format and store it in an Amazon S3 bucket
- H. Import the data into an Amazon Aurora MySQL DB instance
- I. Instruct the data analysts to query the data directly from the DB instance
- J. Send the sensor data to an Amazon Kinesis Data Analytics application to convert the data to Apache Parquet format and store it in an Amazon S3 bucket
- K. Instruct the data analysts to query the data by using Amazon Athena.

Answer: B

NEW QUESTION 225

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

SAP-C02 Practice Exam Features:

- * SAP-C02 Questions and Answers Updated Frequently
- * SAP-C02 Practice Questions Verified by Expert Senior Certified Staff
- * SAP-C02 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * SAP-C02 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The SAP-C02 Practice Test Here](#)