

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional

<https://www.2passeasy.com/dumps/AWS-Certified-Solutions-Architect-Professional/>



NEW QUESTION 1

- (Exam Topic 2)

An external audit of a company's serverless application reveals IAM policies that grant too many permissions. These policies are attached to the company's AWS Lambda execution roles. Hundreds of the company's Lambda functions have broad access permissions, such as full access to Amazon S3 buckets and Amazon DynamoDB tables. The company wants each function to have only the minimum permissions that the function needs to complete its task.

A solutions architect must determine which permissions each Lambda function needs.

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

- A. Set up Amazon CodeGuru to profile the Lambda functions and search for AWS API call
- B. Create an inventory of the required API calls and resources for each Lambda function
- C. Create new IAM access policies for each Lambda function
- D. Review the new policies to ensure that they meet the company's business requirements.
- E. Turn on AWS CloudTrail logging for the AWS account
- F. Use AWS Identity and Access Management Access Analyzer to generate IAM access policies based on the activity recorded in the CloudTrail logs
- G. Review the generated policies to ensure that they meet the company's business requirements.
- H. Turn on AWS CloudTrail logging for the AWS account
- I. Create a script to parse the CloudTrail log, search for AWS API calls by Lambda execution role, and create a summary report
- J. Review the report
- K. Create IAM access policies that provide more restrictive permissions for each Lambda function.
- L. Turn on AWS CloudTrail logging for the AWS account
- M. Export the CloudTrail logs to Amazon S3. Use Amazon EMR to process the CloudTrail logs in Amazon S3 and produce a report of API calls and resources used by each execution role
- N. Create a new IAM access policy for each role
- O. Export the generated roles to an S3 bucket
- P. Review the generated policies to ensure that they meet the company's business requirements.

Answer: B

Explanation:

IAM Access Analyzer helps you identify the resources in your organization and accounts, such as Amazon S3 buckets or IAM roles, shared with an external entity. This lets you identify unintended access to your resources and data, which is a security risk. IAM Access Analyzer identifies resources shared with external principals by using logic-based reasoning to analyze the resource-based policies in your AWS environment.

<https://docs.aws.amazon.com/IAM/latest/UserGuide/what-is-access-analyzer.html>

NEW QUESTION 2

- (Exam Topic 2)

A company has deployed its database on an Amazon RDS for MySQL DB instance in the us-east-1 Region. The company needs to make its data available to customers in Europe. The customers in Europe must have access to the same data as customers in the United States (US) and will not tolerate high application latency or stale data. The customers in Europe and the customers in the US need to write to the database. Both groups of customers need to see updates from the other group in real time.

Which solution will meet these requirements?

- A. Create an Amazon Aurora MySQL replica of the RDS for MySQL DB instance
- B. Pause application writes to the RDS DB instance
- C. Promote the Aurora Replica to a standalone DB instance
- D. Reconfigure the application to use the Aurora database and resume writes
- E. Add eu-west-1 as a secondary Region to the DB instance
- F. Enable write forwarding on the DB instance
- G. Deploy the application in eu-west-1. Configure the application to use the Aurora MySQL endpoint in eu-west-1.
- H. Add a cross-Region replica in eu-west-1 for the RDS for MySQL DB instance
- I. Configure the replica to replicate write queries back to the primary DB instance
- J. Deploy the application in eu-west-1. Configure the application to use the RDS for MySQL endpoint in eu-west-1.
- K. Copy the most recent snapshot from the RDS for MySQL DB instance to eu-west-1. Create a new RDS for MySQL DB instance in eu-west-1 from the snapshot
- L. Configure MySQL logical replication from us-east-1 to eu-west-1. Enable write forwarding on the DB instance
- M. Deploy the application in eu-west-1. Configure the application to use the RDS for MySQL endpoint in eu-west-1.
- N. Convert the RDS for MySQL DB instance to an Amazon Aurora MySQL DB instance
- O. Add eu-west-1 as a secondary Region to the DB instance
- P. Enable write forwarding on the DB instance
- Q. Deploy the application in eu-west-1. Configure the application to use the Aurora MySQL endpoint in eu-west-1.

Answer: D

Explanation:

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

- Amplify Studio: A visual interface that enables you to build and deploy a full-stack app quickly, including frontend UI and backend.
- Amplify CLI: A local toolchain that enables you to configure and manage an app backend with just a few commands.
- Amplify Libraries: Open-source client libraries that enable you to build cloud-powered mobile and web apps.
- Amplify UI Components: Open-source design system with cloud-connected components for building feature-rich apps fast.
- Amplify Hosting: Fully managed CI/CD and hosting for fast, secure, and reliable static and server-side rendered apps.

By using AWS Amplify to create a static website for uploads of media files, the company can leverage Amplify Studio to visually build a pixel-perfect UI and connect it to a cloud backend in clicks. By using Amplify Hosting to serve the website through Amazon CloudFront, the company can easily deploy its web app or website to the fast, secure, and reliable AWS content delivery network (CDN), with hundreds of points of presence globally. By using Amazon S3 to store the uploaded media files, the company can benefit from a highly scalable, durable, and cost-effective object storage service that can handle any amount of data. By

using Amazon Cognito to authenticate users, the company can add user sign-up, sign-in, and access control to its web app with a fully managed service that scales to support millions of users³.

The other options are not correct because:

- Using AWS Application Migration Service to migrate the application server to Amazon EC2 instances would not refactor the application or accelerate development. AWS Application Migration Service (AWS MGN) is a service that enables you to migrate physical servers, virtual machines (VMs), or cloud servers from any source infrastructure to AWS without requiring agents or specialized tools. However, this would not address the challenges of overutilization and data uploads failures. It would also not reduce operational overhead or costs compared to a serverless architecture.
- Creating a static website for uploads of media files and using AWS AppSync to create an API would not be as simple or fast as using AWS Amplify. AWS AppSync is a service that enables you to create flexible APIs for securely accessing, manipulating, and combining data from one or more data sources. However, this would require more configuration and management than using Amplify Studio and Amplify Hosting. It would also not provide authentication features like Amazon Cognito.
- Setting up AWS IAM Identity Center (AWS Single Sign-On) to give users the ability to sign in to the application would not be as suitable as using Amazon Cognito. AWS Single Sign-On (AWS SSO) is a service that enables you to centrally manage SSO access and user permissions across multiple AWS accounts and business applications. However, this service is designed for enterprise customers who need to manage access for employees or partners across multiple resources. It is not intended for authenticating end users of web or mobile apps.

References:

- <https://aws.amazon.com/amplify/>
- <https://aws.amazon.com/s3/>
- <https://aws.amazon.com/cognito/>
- <https://aws.amazon.com/mgn/>
- <https://aws.amazon.com/appsync/>
- <https://aws.amazon.com/single-sign-on/>

NEW QUESTION 3

- (Exam Topic 2)

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 4

- (Exam Topic 2)

A company has VPC flow logs enabled for its NAT gateway. The company is seeing Action = ACCEPT for inbound traffic that comes from public IP address 198.51.100.2 destined for a private Amazon EC2 instance.

A solutions architect must determine whether the traffic represents unsolicited inbound connections from the internet. The first two octets of the VPC CIDR block are 203.0.

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

- A. Open the AWS CloudTrail console
- B. Select the log group that contains the NAT gateway's elastic network interface and the private instance's elastic network interface
- C. Run a query to filter with the destination address set as "like 203.0" and the source address set as "like 198.51.100.2". Run the stats command to filter the sum of bytes transferred by the source address and the destination address.
- D. Open the Amazon CloudWatch console
- E. Select the log group that contains the NAT gateway's elastic network interface and the private instance's elastic network interface
- F. Run a query to filter with the destination address set as "like 203.0" and the source address set as "like 198.51.100.2". Run the stats command to filter the sum of bytes transferred by the source address and the destination address.
- G. Open the AWS CloudTrail console
- H. Select the log group that contains the NAT gateway's elastic network interface and the private instance's elastic network interface
- I. Run a query to filter with the destination address set as "like 198.51.100.2" and the source address set as "like 203.0". Run the stats command to filter the sum of bytes transferred by the source address and the destination address.
- J. Open the Amazon CloudWatch console
- K. Select the log group that contains the NAT gateway's elastic network interface and the private instance's elastic network interface
- L. Run a query to filter with the destination address set as "like 198.51.100.2" and the source address set as "like 203.0". Run the stats command to filter the sum of bytes transferred by the source address and the destination address.

Answer: D

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/vpc-analyze-inbound-traffic-nat-gateway/> by Cloudxie says "select appropriate log"

NEW QUESTION 5

- (Exam Topic 2)

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

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

Answer: A

Explanation:

<https://aws.amazon.com/blogs/compute/implementing-canary-deployments-of-aws-lambda-functions-with-alias>

NEW QUESTION 6

- (Exam Topic 2)

A company needs to optimize the cost of an AWS environment that contains multiple accounts in an organization in AWS Organizations. The company conducted cost optimization activities 3 years ago and purchased Amazon EC2 Standard Reserved Instances that recently expired. The company needs EC2 instances for 3 more years. Additionally, the company has deployed a new serverless workload. Which strategy will provide the company with the MOST cost savings?

- A. Purchase the same Reserved Instances for an additional 3-year term with All Upfront payment
- B. Purchase a 3-year Compute Savings Plan with All Upfront payment in the management account to cover any additional compute costs.
- C. Purchase a 1-year Compute Savings Plan with No Upfront payment in each member account
- D. Use the Savings Plans recommendations in the AWS Cost Management console to choose the Compute Savings Plan.
- E. Purchase a 3-year EC2 Instance Savings Plan with No Upfront payment in the management account to cover EC2 costs in each AWS Region
- F. Purchase a 3-year Compute Savings Plan with No Upfront payment in the management account to cover any additional compute costs.
- G. Purchase a 3-year EC2 Instance Savings Plan with All Upfront payment in each member account
- H. Use the Savings Plans recommendations in the AWS Cost Management console to choose the EC2 Instance Savings Plan.

Answer: A

Explanation:

The company should purchase the same Reserved Instances for an additional 3-year term with All Upfront payment. The company should purchase a 3-year Compute Savings Plan with All Upfront payment in the management account to cover any additional compute costs. This solution will provide the company with the most cost savings because Reserved Instances and Savings Plans are both pricing models that offer significant discounts compared to On-Demand pricing. Reserved Instances are commitments to use a specific instance type and size in a single Region for a one- or three-year term. You can choose between three payment options:

No Upfront, Partial Upfront, or All Upfront. The more you pay upfront, the greater the discount. Savings Plans are flexible pricing models that offer low prices on EC2 instances, Fargate, and Lambda usage, in exchange for a commitment to a consistent amount of usage (measured in \$/hour) for a one- or three-year term. You can choose between two types of Savings Plans: Compute Savings Plans and EC2 Instance Savings Plans. Compute Savings Plans apply to any EC2 instance regardless of Region, instance family, operating system, or tenancy, including those that are part of EMR, ECS, or EKS clusters, or launched by Fargate or Lambda. EC2 Instance Savings Plans apply to a specific instance family within a Region and provide the most savings. By purchasing the same Reserved Instances for an additional 3-year term with All Upfront payment, the company can lock in the lowest possible price for its EC2 instances that run continuously for 3 years. By purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account, the company can benefit from additional discounts on any other compute usage across its member accounts.

The other options are not correct because:

- Purchasing a 1-year Compute Savings Plan with No Upfront payment in each member account would not provide as much cost savings as purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account. A 1-year term offers lower discounts than a 3-year term, and a No Upfront payment option offers lower discounts than an All Upfront payment option. Also, purchasing a Savings Plan in each member account would not allow the company to share the benefits of unused Savings Plan discounts across its organization.
- Purchasing a 3-year EC2 Instance Savings Plan with No Upfront payment in the management account to cover EC2 costs in each AWS Region would not provide as much cost savings as purchasing Reserved Instances for an additional 3-year term with All Upfront payment. An EC2 Instance Savings Plan offers lower discounts than Reserved Instances for the same instance family and Region. Also, a No Upfront payment option offers lower discounts than an All Upfront payment option.
- Purchasing a 3-year EC2 Instance Savings Plan with All Upfront payment in each member account would not provide as much flexibility or cost savings as purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account. An EC2 Instance Savings Plan applies only to a specific instance family within a Region and does not cover Fargate or Lambda usage. Also, purchasing a Savings Plan in each member account would not allow the company to share the benefits of unused Savings Plan discounts across its organization.

References:

- <https://aws.amazon.com/ec2/pricing/reserved-instances/>
- <https://aws.amazon.com/savingsplans/>

NEW QUESTION 7

- (Exam Topic 2)

A solutions architect needs to improve an application that is hosted in the AWS Cloud. The application uses an Amazon Aurora MySQL DB instance that is experiencing overloaded connections. Most of the application's operations insert records into the database. The application currently stores credentials in a text-based configuration file.

The solutions architect needs to implement a solution so that the application can handle the current connection load. The solution must keep the credentials secure and must provide the ability to rotate the credentials automatically on a regular basis.

Which solution will meet these requirements?

- A. Deploy an Amazon RDS Proxy layer in front of the DB instance
- B. Store the connection credentials as a secret in AWS Secrets Manager.

- C. Deploy an Amazon RDS Proxy layer in front of the DB instance
- D. Store the connection credentials in AWS Systems Manager Parameter Store.
- E. Create an Aurora Replica
- F. Store the connection credentials as a secret in AWS Secrets Manager.
- G. Create an Aurora Replica
- H. Store the connection credentials in AWS Systems Manager Parameter Store.

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/rds-proxy.html>

NEW QUESTION 8

- (Exam Topic 2)

A company recently started hosting new application workloads in the AWS Cloud. The company is using Amazon EC2 instances, Amazon Elastic File System (Amazon EFS) file systems, and Amazon RDS DB instances.

To meet regulatory and business requirements, the company must make the following changes for data backups:

- * Backups must be retained based on custom daily, weekly, and monthly requirements.
- * Backups must be replicated to at least one other AWS Region immediately after capture.
- * The backup solution must provide a single source of backup status across the AWS environment.
- * The backup solution must send immediate notifications upon failure of any resource backup.

Which combination of steps will meet this requirement with the LEAST amount of operational overhead? (Select THREE.)

- A. Create an AWS Backup plan with a backup rule for each of the retention requirements.
- B. Configure an AWS backup plan to copy backups to another Region.
- C. Create an AWS Lambda function to replicate backups to another Region and send notification if a failure occurs.
- D. Add an Amazon Simple Notification Service (Amazon SNS) topic to the backup plan to send a notification for finished jobs that have any status except BACKUP- JOB- COMPLETED.
- E. Create an Amazon Data Lifecycle Manager (Amazon DLM) snapshot lifecycle policy for each of the retention requirements.
- F. Set up RDS snapshots on each database.

Answer: ABD

Explanation:

Cross region with AWS Backup:

<https://docs.aws.amazon.com/aws-backup/latest/devguide/cross-region-backup.html>

NEW QUESTION 9

- (Exam Topic 2)

A company needs to build a disaster recovery (DR) solution for its ecommerce website. The web application is hosted on a fleet of t3.large Amazon EC2 instances and uses an Amazon RDS for MySQL DB instance. The EC2 instances are in an Auto Scaling group that extends across multiple Availability Zones.

In the event of a disaster, the web application must fail over to the secondary environment with an RPO of 30 seconds and an RTO of 10 minutes.

Which solution will meet these requirements MOST cost-effectively?

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

Answer: B

Explanation:

The company should use infrastructure as code (IaC) to provision the new infrastructure in the DR Region. The company should create a cross-Region read replica for the DB instance. The company should set up AWS Elastic Disaster Recovery to continuously replicate the EC2 instances to the DR Region. The company should run the EC2 instances at the minimum capacity in the DR Region. The company should use an Amazon Route 53 failover routing policy to automatically fail over to the DR Region in the event of a disaster. The company should increase the desired capacity of the Auto Scaling group. This solution will meet the requirements most cost-effectively because AWS Elastic Disaster Recovery (AWS DRS) is a service that minimizes downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications using affordable storage, minimal compute, and point-in-time recovery. AWS DRS enables RPOs of seconds and RTOs of minutes. AWS DRS continuously replicates data from the source servers to a staging area subnet in the DR Region, where it uses low-cost storage and minimal compute resources to maintain ongoing replication. In the event of a disaster, AWS DRS automatically converts the servers to boot and run natively on AWS and launches recovery instances on AWS within minutes. By using AWS DRS, the company can save costs by removing idle recovery site resources and paying for the full disaster recovery site only when needed. By creating a cross-Region read replica for the DB instance, the company can have a standby copy of its primary database in a different AWS Region. By using infrastructure as code (IaC), the company can provision the new infrastructure in the

DR Region in an automated and consistent way⁴. By using an Amazon Route 53 failover routing policy, the company can route traffic to a resource that is healthy or to another resource when the first resource becomes unavailable.

The other options are not correct because:

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

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

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

References:

- <https://aws.amazon.com/disaster-recovery/>
- <https://docs.aws.amazon.com/drs/latest/userguide/what-is-drs.html>
- <https://aws.amazon.com/cloudformation/>
- <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover.html>
- <https://aws.amazon.com/backup/>
- <https://docs.aws.amazon.com/redshift/latest/mgmt/data-api.html>
- <https://aws.amazon.com/data-exchange/>
- <https://docs.aws.amazon.com/redshift/latest/dg/datashare-overview.html>

NEW QUESTION 10

- (Exam Topic 2)

A company is designing a new website that hosts static content. The website will give users the ability to upload and download large files. According to company requirements, all data must be encrypted in transit and at rest. A solutions architect is building the solution by using Amazon S3 and Amazon CloudFront. Which combination of steps will meet the encryption requirements? (Select THREE.)

- A. Turn on S3 server-side encryption for the S3 bucket that the web application uses.
- B. Add a policy attribute of "aws:SecureTransport": "true" for read and write operations in the S3 ACLs.
- C. Create a bucket policy that denies any unencrypted operations in the S3 bucket that the web application uses.
- D. Configure encryption at rest on CloudFront by using server-side encryption with AWS KMS keys (SSE-KMS).
- E. Configure redirection of HTTP requests to HTTPS requests in CloudFront.
- F. Use the RequireSSL option in the creation of presigned URLs for the S3 bucket that the web application uses.

Answer: ACE

Explanation:

Turning on S3 server-side encryption for the S3 bucket that the web application uses will enable encrypting the data at rest using Amazon S3 managed keys (SSE-S3)¹. Creating a bucket policy that denies any unencrypted operations in the S3 bucket that the web application uses will enable enforcing encryption for all requests to the bucket². Configuring redirection of HTTP requests to HTTPS requests in CloudFront will enable encrypting the data in transit using SSL/TLS³.

NEW QUESTION 10

- (Exam Topic 2)

A company is storing sensitive data in an Amazon S3 bucket. The company must log all activities for objects in the S3 bucket and must keep the logs for 5 years. The company's security team also must receive an email notification every time there is an attempt to delete data in the S3 bucket. Which combination of steps will meet these requirements MOST cost-effectively? (Select THREE.)

- A. Configure AWS CloudTrail to log S3 data events.
- B. Configure S3 server access logging for the S3 bucket.
- C. Configure Amazon S3 to send object deletion events to Amazon Simple Email Service (Amazon SES).
- D. Configure Amazon S3 to send object deletion events to an Amazon EventBridge event bus that publishes to an Amazon Simple Notification Service (Amazon SNS) topic.
- E. Configure Amazon S3 to send the logs to Amazon Timestream with data storage tiering.
- F. Configure a new S3 bucket to store the logs with an S3 Lifecycle policy.

Answer: ADF

Explanation:

Configuring AWS CloudTrail to log S3 data events will enable logging all activities for objects in the S3 bucket¹. Data events are object-level API operations such as GetObject, DeleteObject, and PutObject¹. Configuring Amazon S3 to send object deletion events to an Amazon EventBridge event bus that publishes to an Amazon Simple Notification Service (Amazon SNS) topic will enable sending email notifications every time there is an attempt to delete data in the S3 bucket². EventBridge can route events from S3 to SNS, which can send emails to subscribers². Configuring a new S3 bucket to store the logs with an S3 Lifecycle policy will enable keeping the logs for 5 years in a cost-effective way³. A lifecycle policy can transition the logs to a cheaper storage class such as Glacier or delete them after a specified period of time³.

NEW QUESTION 14

- (Exam Topic 2)

A company uses AWS Organizations to manage more than 1,000 AWS accounts. The company has created a new developer organization. There are 540 developer member accounts that must be moved to the new developer organization. All accounts are set up with all the required information so that each account can be operated as a standalone account.

Which combination of steps should a solutions architect take to move all of the developer accounts to the new developer organization? (Select THREE.)

- A. Call the MoveAccount operation in the Organizations API from the old organization's management account to migrate the developer accounts to the new developer organization.
- B. From the management account, remove each developer account from the old organization using the RemoveAccountFromOrganization operation in the Organizations API.
- C. From each developer account, remove the account from the old organization using the RemoveAccountFromOrganization operation in the Organizations API.
- D. Sign in to the new developer organization's management account and create a placeholder member account that acts as a target for the developer account migration.
- E. Call the InviteAccountToOrganization operation in the Organizations API from the new developer organization's management account to send invitations to the developer accounts.
- F. Have each developer sign in to their account and confirm to join the new developer organization.

Answer: BEF

Explanation:

"This operation can be called only from the organization's management account. Member accounts can remove themselves with LeaveOrganization instead."
https://docs.aws.amazon.com/organizations/latest/APIReference/API_RemoveAccountFromOrganization.html

NEW QUESTION 15

- (Exam Topic 2)

A company is running an application in the AWS Cloud. The application collects and stores a large amount of unstructured data in an Amazon S3 bucket. The S3 bucket contains several terabytes of data and uses the S3 Standard storage class. The data increases in size by several gigabytes every day. The company needs to query and analyze the data. The company does not access data that is more than 1 year old. However, the company must retain all the data indefinitely for compliance reasons.

Which solution will meet these requirements MOST cost-effectively?

- A. Use S3 Select to query the data
- B. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Glacier Deep Archive.
- C. Use Amazon Redshift Spectrum to query the data
- D. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Glacier Deep Archive.
- E. Use an AWS Glue Data Catalog and Amazon Athena to query the data
- F. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Glacier Deep Archive.
- G. Use Amazon Redshift Spectrum to query the data
- H. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Intelligent-Tiering.

Answer: C

Explanation:

Generally, unstructured data should be converted structured data before querying them. AWS Glue can do that.
<https://docs.aws.amazon.com/glue/latest/dg/schema-relationalize.html> <https://docs.aws.amazon.com/athena/latest/ug/glue-athena.html>

NEW QUESTION 20

- (Exam Topic 2)

A solutions architect at a large company needs to set up network security for outbound traffic to the internet from all AWS accounts within an organization in AWS Organizations. The organization has more than 100 AWS accounts, and the accounts route to each other by using a centralized AWS Transit Gateway. Each account has both an internet gateway and a NAT gateway for outbound traffic to the internet. The company deploys resources only into a single AWS Region. The company needs the ability to add centrally managed rule-based filtering on all outbound traffic to the internet for all AWS accounts in the organization. The peak load of outbound traffic will not exceed 25 Gbps in each Availability Zone.

Which solution meets these requirements?

- A. Create a new VPC for outbound traffic to the internet
- B. Connect the existing transit gateway to the new VPC
- C. Configure a new NAT gateway
- D. Create an Auto Scaling group of Amazon EC2 instances that run an open-source internet proxy for rule-based filtering across all Availability Zones in the Region
- E. Modify all default routes to point to the proxy's Auto Scaling group.
- F. Create a new VPC for outbound traffic to the internet
- G. Connect the existing transit gateway to the new VPC
- H. Configure a new NAT gateway
- I. Use an AWS Network Firewall firewall for rule-based filtering
- J. Create Network Firewall endpoints in each Availability Zone
- K. Modify all default routes to point to the Network Firewall endpoints.
- L. Create an AWS Network Firewall firewall for rule-based filtering in each AWS account
- M. Modify all default routes to point to the Network Firewall firewalls in each account.
- N. In each AWS account, create an Auto Scaling group of network-optimized Amazon EC2 instances that run an open-source internet proxy for rule-based filtering
- O. Modify all default routes to point to the proxy's Auto Scaling group.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/networking-and-content-delivery/deployment-models-for-aws-network-firewall/>

NEW QUESTION 21

- (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 VPC
- B. Configure the required routing to allow access to the internet.

- C. Create a transit gateway, and share it with the existing AWS account
- D. Attach existing VPCs to the transit gateway Configure the required routing to allow access to the internet.
- E. Create a transit gateway in every account
- F. Attach the NAT gateway to the transit gateway
- G. Configure the required routing to allow access to the internet.
- H. Create an AWS PrivateLink connection between the egress VPC and the spoke VPC
- I. Configure the required routing to allow access to the internet

Answer: B

Explanation:

<https://d1.awsstatic.com/architecture-diagrams/ArchitectureDiagrams/NAT-gateway-centralized-egress-ra.pdf?d>

NEW QUESTION 22

- (Exam Topic 2)

A company is using AWS Organizations to manage multiple AWS accounts. For security purposes, the company requires the creation of an Amazon Simple Notification Service (Amazon SNS) topic that enables integration with a third-party alerting system in all the Organizations member accounts.

A solutions architect used an AWS CloudFormation template to create the SNS topic and stack sets to automate the deployment of Cloud Formation stacks. Trusted access has been enabled in Organizations.

What should the solutions architect do to deploy the CloudFormation StackSets in all AWS accounts?

- A. Create a stack set in the Organizations member account
- B. Use service-managed permission
- C. Set deployment options to deploy to an organization
- D. Use CloudFormation StackSets drift detection.
- E. Create stacks in the Organizations member account
- F. Use self-service permission
- G. Set deployment options to deploy to an organization
- H. Enable the CloudFormation StackSets automatic deployment.
- I. Create a stack set in the Organizations management account
- J. Use service-managed permission
- K. Set deployment options to deploy to the organization
- L. Enable CloudFormation StackSets automatic deployment.
- M. Create stacks in the Organizations management account
- N. Use service-managed permission
- O. Set deployment options to deploy to the organization
- P. Enable CloudFormation StackSets drift detection.

Answer: C

Explanation:

[https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-orgs-manage-auto-deployment.h](https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-orgs-manage-auto-deployment.html)

NEW QUESTION 23

- (Exam Topic 2)

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 to 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 Attach an extra elastic network interface to each EC2 instance.

Answer: A

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html#placement-groups-cluster>

NEW QUESTION 24

- (Exam Topic 2)

A company is running a critical stateful web application on two Linux Amazon EC2 instances behind an Application Load Balancer (ALB) with an Amazon RDS for MySQL database. The company hosts the DNS records for the application in Amazon Route 53. A solutions architect must recommend a solution to improve the resiliency of the application.

The solution must meet the following objectives:

- Application tier RPO of 2 minutes. RTO of 30 minutes
- Database tier RPO of 5 minutes RTO of 30 minutes

The company does not want to make significant changes to the existing application architecture. The company must ensure optimal latency after a failover. Which solution will meet these requirements?

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

second Region Configure an Amazon CloudFront distribution in front of the ALB Update DNS records to point to CloudFront
D. Configure the EC2 instances to use Amazon Data Lifecycle Manager (Amazon DLM) to take snapshots of the EBS volumes Create a cross-Region read replica for the RDS DB instance Create an ALB in a second AWS Region Create an AWS Global Accelerator endpoint and associate the endpoint with the ALBs

Answer: B

Explanation:

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

NEW QUESTION 28

- (Exam Topic 2)

A company wants to optimize AWS data-transfer costs and compute costs across developer accounts within the company's organization in AWS Organizations Developers can configure VPCs and launch Amazon EC2 instances in a single AWS Region The EC2 instances retrieve approximately 1 TB of data each day from Amazon S3

The developer activity leads to excessive monthly data-transfer charges and NAT gateway processing charges between EC2 instances and S3 buckets, along with high compute costs The company wants to proactively enforce approved architectural patterns for any EC2 instance and VPC infrastructure that developers deploy within the AWS accounts The company does not want this enforcement to negatively affect the speed at which the developers can perform their tasks Which solution will meet these requirements MOST cost-effectively?

- A. Create SCPs to prevent developers from launching unapproved EC2 instance types Provide the developers with an AWS CloudFormation template to deploy an approved VPC configuration with S3 interface endpoints Scope the developers' IAM permissions so that the developers can launch VPC resources only with CloudFormation
- B. Create a daily forecasted budget with AWS Budgets to monitor EC2 compute costs and S3 data-transfer costs across the developer accounts When the forecasted cost is 75% of the actual budget cost, send an alert to the developer teams If the actual budget cost is 100%. create a budget action to terminate the developers' EC2 instances and VPC infrastructure
- C. Create an AWS Service Catalog portfolio that users can use to create an approved VPC configuration with S3 gateway endpoints and approved EC2 instances Share the portfolio with the developer accounts Configure an AWS Service Catalog launch constraint to use an approved IAM role Scope the developers' IAM permissions to allow access only to AWS Service Catalog
- D. Create and deploy AWS Config rules to monitor the compliance of EC2 and VPC resources in the developer AWS accounts If developers launch unapproved EC2 instances or if developers create VPCs without S3 gateway endpoints perform a remediation action to terminate the unapproved resources

Answer: C

Explanation:

This solution allows developers to quickly launch resources using pre-approved configurations and instance types, while also ensuring that the resources launched comply with the company's architectural patterns. This can help reduce data transfer and compute costs associated with the resources. Using AWS Service Catalog also allows the company to control access to the approved configurations and resources through the use of IAM roles, while also allowing developers to quickly provision resources without negatively affecting their ability to perform their tasks.

Reference:

AWS Service Catalog: <https://aws.amazon.com/service-catalog/> AWS Service Catalog Constraints:

<https://docs.aws.amazon.com/servicecatalog/latest/adminguide/constraints.html>

AWS Service Catalog Launch Constraints: <https://docs.aws.amazon.com/servicecatalog/latest/adminguide/launch-constraints.html>

NEW QUESTION 33

- (Exam Topic 2)

A company hosts a blog post application on AWS using Amazon API Gateway, Amazon DynamoDB, and AWS Lambda. The application currently does not use API keys to authorize requests. The API model is as follows: GET/posts/[postid] to get post details GET/users[userid] to get user details GET/comments/[commentid] to get comments details

The company has noticed users are actively discussing topics in the comments section, and the company wants to increase user engagement by marking the comments appears in real time.

Which design should be used to reduce comment latency and improve user experience?

- A. Use edge-optimized API with Amazon CloudFront to cache API responses.
- B. Modify the blog application code to request GET comment[commented] every 10 seconds.
- C. Use AWS AppSync and leverage WebSockets to deliver comments.
- D. Change the concurrency limit of the Lambda functions to lower the API response time.

Answer: C

Explanation:

<https://docs.aws.amazon.com/appsync/latest/devguide/graphql-overview.html>

AWS AppSync is a fully managed GraphQL service that allows applications to securely access, manipulate, and receive data as well as real-time updates from multiple data sources¹. AWS AppSync supports GraphQL subscriptions to perform real-time operations and can push data to clients that choose to listen to specific events from the backend¹. AWS AppSync uses WebSockets to establish and maintain a secure connection between the clients and the API endpoint². Therefore, using AWS AppSync and leveraging WebSockets is a suitable design to reduce comment latency and improve user experience.

NEW QUESTION 38

- (Exam Topic 2)

A company uses an AWS CodeCommit repository The company must store a backup copy of the data that is in the repository in a second AWS Region Which solution will meet these requirements?

- A. Configure AWS Elastic Disaster Recovery to replicate the CodeCommit repository data to the second Region
- B. Use AWS Backup to back up the CodeCommit repository on an hourly schedule Create a cross-Region copy in the second Region
- C. Create an Amazon EventBridge rule to invoke AWS CodeBuild when the company pushes code to the repository Use CodeBuild to clone the repository Create a zip file of the content Copy the file to an S3 bucket in the second Region
- D. Create an AWS Step Functions workflow on an hourly schedule to take a snapshot of the CodeCommit repository Configure the workflow to copy the snapshot to an S3 bucket in the second Region

Answer: B

Explanation:

AWS Backup is a fully managed service that makes it easy to centralize and automate the creation, retention, and restoration of backups across AWS services. It provides a way to schedule automatic backups for CodeCommit repositories on an hourly basis. Additionally, it also supports cross-Region replication, which allows you to copy the backups to a second Region for disaster recovery.

By using AWS Backup, the company can set up an automatic and regular backup schedule for the CodeCommit repository, ensuring that the data is regularly backed up and stored in a second Region. This can provide a way to recover quickly from any disaster event that might occur.

Reference:

AWS Backup documentation: <https://aws.amazon.com/backup/> AWS Backup for AWS CodeCommit documentation:

<https://aws.amazon.com/about-aws/whats-new/2020/07/aws-backup-now-supports-aws-codecommit-repositorie>

NEW QUESTION 40

- (Exam Topic 2)

A company has several AWS accounts. A development team is building an automation framework for cloud governance and remediation processes. The automation framework uses AWS Lambda functions in a centralized account. A solutions architect must implement a least privilege permissions policy that allows the Lambda functions to run in each of the company's AWS accounts.

Which combination of steps will meet these requirements? (Choose two.)

- A. In the centralized account, create an IAM role that has the Lambda service as a trusted entity
- B. Add an inline policy to assume the roles of the other AWS accounts.
- C. In the other AWS accounts, create an IAM role that has minimal permission
- D. Add the centralized account's Lambda IAM role as a trusted entity.
- E. In the centralized account, create an IAM role that has roles of the other accounts as trusted entities. Provide minimal permissions.
- F. In the other AWS accounts, create an IAM role that has permissions to assume the role of the centralized account
- G. Add the Lambda service as a trusted entity.
- H. In the other AWS accounts, create an IAM role that has minimal permission
- I. Add the Lambda service as a trusted entity.

Answer: AB

Explanation:

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

NEW QUESTION 41

- (Exam Topic 2)

A company's public API runs as tasks on Amazon Elastic Container Service (Amazon ECS). The tasks run on AWS Fargate behind an Application Load Balancer (ALB) and are configured with Service Auto Scaling for the tasks based on CPU utilization. This service has been running well for several months.

Recently, API performance slowed down and made the application unusable. The company discovered that a significant number of SQL injection attacks had occurred against the API and that the API service had scaled to its maximum amount.

A solutions architect needs to implement a solution that prevents SQL injection attacks from reaching the ECS API service. The solution must allow legitimate traffic through and must maximize operational efficiency. Which solution meets these requirements?

- A. Create a new AWS WAF web ACL to monitor the HTTP requests and HTTPS requests that are forwarded to the ALB in front of the ECS tasks.
- B. Create a new AWS WAF Bot Control implementation
- C. Add a rule in the AWS WAF Bot Control managed rule group to monitor traffic and allow only legitimate traffic to the ALB in front of the ECS tasks.
- D. Create a new AWS WAF web ACL
- E. Add a new rule that blocks requests that match the SQL database rule group
- F. Set the web ACL to allow all other traffic that does not match those rules
- G. Attach the web ACL to the ALB in front of the ECS tasks.
- H. Create a new AWS WAF web ACL
- I. Create a new empty IP set in AWS WAF
- J. Add a new rule to the web ACL to block requests that originate from IP addresses in the new IP set
- K. Create an AWS Lambda function that scrapes the API logs for IP addresses that send SQL injection attacks, and add those IP addresses to the IP set
- L. Attach the web ACL to the ALB in front of the ECS tasks.

Answer: C

Explanation:

The company should create a new AWS WAF web ACL. The company should add a new rule that blocks requests that match the SQL database rule group. The company should set the web ACL to allow all other traffic that does not match those rules. The company should attach the web ACL to the ALB in front of the ECS tasks. This solution will meet the requirements because AWS WAF is a web application firewall that lets you monitor and control web requests that are forwarded to your web applications. You can use AWS WAF to define customizable web security rules that control which traffic can access your web applications and which traffic should be blocked¹. By creating a new AWS WAF web ACL, the company can create a collection of rules that define the conditions for allowing or blocking web requests. By adding a new rule that blocks requests that match the SQL database rule group, the company can prevent SQL injection attacks from reaching the ECS API service. The SQL database rule group is a managed rule group provided by AWS that contains rules to protect against common SQL injection attack patterns². By setting the web ACL to allow all other traffic that does not match those rules, the company can ensure that legitimate traffic can access the API service. By attaching the web ACL to the ALB in front of the ECS tasks, the company can apply the web security rules to all requests that are forwarded by the load balancer.

The other options are not correct because:

- Creating a new AWS WAF Bot Control implementation would not prevent SQL injection attacks from reaching the ECS API service. AWS WAF Bot Control is a feature that gives you visibility and control over common and pervasive bot traffic that can consume excess resources, skew metrics, cause downtime, or perform other undesired activities. However, it does not protect against SQL injection attacks, which are malicious attempts to execute unauthorized SQL statements against your database³.
- Creating a new AWS WAF web ACL to monitor the HTTP requests and HTTPS requests that are forwarded to the ALB in front of the ECS tasks would not prevent SQL injection attacks from reaching the ECS API service. Monitoring mode is a feature that enables you to evaluate how your rules would perform without actually blocking any requests. However, this mode does not provide any protection against attacks, as it only logs and counts requests that match your rules⁴.
- Creating a new AWS WAF web ACL and creating a new empty IP set in AWS WAF would not prevent SQL injection attacks from reaching the ECS API service. An IP set is a feature that enables you to specify a list of IP addresses or CIDR blocks that you want to allow or block based on their source IP address. However, this approach would not be effective or efficient against SQL injection attacks, as it would require constantly updating the IP set with new IP addresses of attackers, and it would not block attackers who use proxies or VPNs.

References:

- > <https://aws.amazon.com/waf/>
- > <https://docs.aws.amazon.com/waf/latest/developerguide/waf-bot-control.html>
- > <https://docs.aws.amazon.com/waf/latest/developerguide/web-acl-monitoring-mode.html>
- > <https://docs.aws.amazon.com/waf/latest/developerguide/waf-ip-sets.html>

NEW QUESTION 44

- (Exam Topic 2)

A company uses a Grafana data visualization solution that runs on a single Amazon EC2 instance to monitor the health of the company's AWS workloads. The company has invested time and effort to create dashboards that the company wants to preserve. The dashboards need to be highly available and cannot be down for longer than 10 minutes. The company needs to minimize ongoing maintenance.

Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: C

Explanation:

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

NEW QUESTION 47

- (Exam Topic 2)

A company runs an intranet application on premises. The company wants to configure a cloud backup of the application. The company has selected AWS Elastic Disaster Recovery for this solution.

The company requires that replication traffic does not travel through the public internet. The application also must not be accessible from the internet. The company does not want this solution to consume all available network bandwidth because other applications require bandwidth.

Which combination of steps will meet these requirements? (Select THREE.)

- A. Create a VPC that has at least two private subnets, two NAT gateways, and a virtual private gateway.
- B. Create a VPC that has at least two public subnets, a virtual private gateway, and an internet gateway.
- C. Create an AWS Site-to-Site VPN connection between the on-premises network and the target AWS network.
- D. Create an AWS Direct Connect connection and a Direct Connect gateway between the on-premises network and the target AWS network.
- E. During configuration of the replication servers, select the option to use private IP addresses for data replication.
- F. During configuration of the launch settings for the target servers, select the option to ensure that the Recovery instance's private IP address matches the source server's private IP address.

Answer: BDE

Explanation:

AWS Elastic Disaster Recovery (AWS DRS) is a service that minimizes downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications using affordable storage, minimal compute, and point-in-time recovery¹. Users can set up AWS DRS on their source servers to initiate secure data replication to a staging area subnet in their AWS account, in the AWS Region they select. Users can then launch recovery instances on AWS within minutes, using the most up-to-date server state or a previous point in time.

To configure a cloud backup of the application with AWS DRS, users need to create a VPC that has at least two public subnets, a virtual private gateway, and an internet gateway. A VPC is a logically isolated section of the AWS Cloud where users can launch AWS resources in a virtual network that they define². A public subnet is a subnet that has a route to an internet gateway³. A virtual private gateway is the VPN concentrator on the Amazon side of the Site-to-Site VPN connection⁴. An internet gateway is a horizontally scaled, redundant, and highly available VPC component that allows communication between instances in the VPC and the internet. Users need to create at least two public subnets for redundancy and high availability. Users need to create a virtual private gateway and attach it to the VPC to enable VPN connectivity between the on-premises network and the target AWS network. Users need to create an internet gateway and attach it to the VPC to enable internet access for the replication servers.

To ensure that replication traffic does not travel through the public internet, users need to create an AWS Direct Connect connection and a Direct Connect gateway between the on-premises network and the target AWS network. AWS Direct Connect is a service that establishes a dedicated network connection from an on-premises network to one or more VPCs. A Direct Connect gateway is a globally available resource that allows users to connect multiple VPCs across different Regions to their on-premises networks using one or more Direct Connect connections. Users need to create an AWS Direct Connect connection between their on-premises network and an AWS Region. Users need to create a Direct Connect gateway and associate it with their VPC and their Direct Connect connection.

To ensure that the application is not accessible from the internet, users need to select the option to use private IP addresses for data replication during configuration of the replication servers. This option configures the replication servers with private IP addresses only, without assigning any public IP addresses or Elastic IP addresses. This way, the replication servers can only communicate with other resources within the VPC or through VPN connections.

Option A is incorrect because creating a VPC that has at least two private subnets, two NAT gateways, and a virtual private gateway is not necessary or cost-effective. A private subnet is a subnet that does not have a route to an internet gateway³. A NAT gateway is a highly available, managed Network Address Translation (NAT) service that enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating connections with those instances. Users do not need to create private subnets or NAT gateways for this use case, as they can use public subnets with private IP addresses for data replication.

Option C is incorrect because creating an AWS Site-to-Site VPN connection between the on-premises network and the target AWS network will not ensure that replication traffic does not travel through the public internet. A Site-to-Site VPN connection consists of two VPN tunnels between an on-premises customer

gateway device and a virtual private gateway in your VPC4. The VPN tunnels are encrypted using IPSec protocols, but they still use public IP addresses for communication. Users need to use AWS Direct Connect instead of Site-to-Site VPN for this use case.

Option F is incorrect because selecting the option to ensure that the Recovery instance's private IP address matches the source server's private IP address during configuration of the launch settings for the target servers will not ensure that the application is not accessible from the internet. This option configures the Recovery instance with an identical private IP address as its source server when launched in drills or recovery mode. However, this option does not prevent assigning public IP addresses or Elastic IP addresses to the Recovery instance. Users need to select the option to use private IP addresses for data replication instead.

NEW QUESTION 52

- (Exam Topic 2)

A company is creating a centralized logging service running on Amazon EC2 that will receive and analyze logs from hundreds of AWS accounts. AWS PrivateLink is being used to provide connectivity between the client services and the logging service.

In each AWS account with a client, an interface endpoint has been created for the logging service and is available. The logging service running on EC2 instances with a Network Load Balancer (NLB) are deployed in different subnets. The clients are unable to submit logs using the VPC endpoint.

Which combination of steps should a solutions architect take to resolve this issue? (Select TWO.)

- A. Check that the NACL is attached to the logging service subnet to allow communications to and from the NLB subnet
- B. Check that the NACL is attached to the NLB subnet to allow communications to and from the logging service subnets running on EC2 instances.
- C. Check that the NACL is attached to the logging service subnets to allow communications to and from the interface endpoint subnet
- D. Check that the NACL is attached to the interface endpoint subnet to allow communications to and from the logging service subnets running on EC2 instances.
- E. Check the security group for the logging service running on the EC2 instances to ensure it allows Ingress from the NLB subnets.
- F. Check the security group for the logging service running on EC2 instances to ensure it allows ingress from the clients.
- G. Check the security group for the NLB to ensure it allows ingress from the interface endpoint subnets.

Answer: AC

NEW QUESTION 53

- (Exam Topic 2)

A company has many separate AWS accounts and uses no central billing or management. Each AWS account hosts services for different departments in the company. The company has a Microsoft Azure Active Directory that is deployed.

A solution architect needs to centralize billing and management of the company's AWS accounts. The company wants to start using identify federation instead of manual user management. The company also wants to use temporary credentials instead of long-lived access keys.

Which combination of steps will meet these requirements? (Select THREE)

- A. Create a new AWS account to serve as a management account
- B. Deploy an organization in AWS Organization
- C. Invite each existing AWS account to join the organization
- D. Ensure that each account accepts the invitation.
- E. Configure each AWS Account's email address to be aws+<account id>@example.com so that account management email messages and invoices are sent to the same place.
- F. Deploy AWS IAM Identity Center (AWS Single Sign-On) in the management account
- G. Connect IAM Identity Center to the Azure Active Directory
- H. Configure IAM Identity Center for automatic synchronization of users and groups.
- I. Deploy an AWS Managed Microsoft AD directory in the management account
- J. Share the directory with all other accounts in the organization by using AWS Resource Access Manager (AWS RAM).
- K. Create AWS IAM Identity Center (AWS Single Sign-On) permission set
- L. Attach the permission sets to the appropriate IAM Identity Center groups and AWS accounts.
- M. Configure AWS Identity and Access Management (IAM) in each AWS account to use AWS Managed Microsoft AD for authentication and authorization.

Answer: ACE

NEW QUESTION 56

- (Exam Topic 2)

A company needs to establish a connection from its on-premises data center to AWS. The company needs to connect all of its VPCs that are located in different AWS Regions with transitive routing capabilities between VPC networks. The company also must reduce network outbound traffic costs, increase bandwidth throughput, and provide a consistent network experience for end users.

Which solution will meet these requirements?

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

Answer: B

Explanation:

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

NEW QUESTION 61

- (Exam Topic 2)

A company is migrating a document processing workload to AWS. The company has updated many applications to natively use the Amazon S3 API to store,

retrieve, and modify documents that a processing server generates at a rate of approximately 5 documents every second. After the document processing is finished, customers can download the documents directly from Amazon S3. During the migration, the company discovered that it could not immediately update the processing server that generates many documents to support the S3 API. The server runs on Linux and requires fast local access to the files that the server generates and modifies. When the server finishes processing, the files must be available to the public for download within 30 minutes. Which solution will meet these requirements with the LEAST amount of effort?

- A. Migrate the application to an AWS Lambda function
- B. Use the AWS SDK for Java to generate, modify, and access the files that the company stores directly in Amazon S3.
- C. Set up an Amazon S3 File Gateway and configure a file share that is linked to the document store. Mount the file share on an Amazon EC2 instance by using NFS
- D. When changes occur in Amazon S3, initiate a RefreshCache API call to update the S3 File Gateway.
- E. Configure Amazon FSx for Lustre with an import and export policy
- F. Link the new file system to an S3 bucket
- G. Install the Lustre client and mount the document store to an Amazon EC2 instance by using NFS.
- H. Configure AWS DataSync to connect to an Amazon EC2 instance
- I. Configure a task to synchronize the generated files to and from Amazon S3.

Answer: C

Explanation:

Amazon FSx for Lustre is a fully managed service that provides cost-effective, high-performance, scalable storage for compute workloads. Powered by Lustre, the world's most popular high-performance file system, FSx for Lustre offers shared storage with sub-ms latencies, up to terabytes per second of throughput, and millions of IOPS. FSx for Lustre file systems can also be linked to Amazon Simple Storage Service (S3) buckets, allowing you to access and process data concurrently from both a high-performance file system and from the S3 API.

NEW QUESTION 64

- (Exam Topic 2)

A company has an application that runs as a ReplicaSet of multiple pods in an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The EKS cluster has nodes in multiple Availability Zones. The application generates many small files that must be accessible across all running instances of the application. The company needs to back up the files and retain the backups for 1 year.

Which solution will meet these requirements while providing the FASTEST storage performance?

- A. Create an Amazon Elastic File System (Amazon EFS) file system and a mount target for each subnet that contains nodes in the EKS cluster
- B. Configure the ReplicaSet to mount the file system
- C. Direct the application to store files in the file system
- D. Configure AWS Backup to back up and retain copies of the data for 1 year.
- E. Create an Amazon Elastic Block Store (Amazon EBS) volume
- F. Enable the EBS Multi-Attach feature. Configure the ReplicaSet to mount the EBS volume
- G. Direct the application to store files in the EBS volume
- H. Configure AWS Backup to back up and retain copies of the data for 1 year.
- I. Create an Amazon S3 bucket
- J. Configure the ReplicaSet to mount the S3 bucket
- K. Direct the application to store files in the S3 bucket
- L. Configure S3 Versioning to retain copies of the data
- M. Configure an S3 Lifecycle policy to delete objects after 1 year.
- N. Configure the ReplicaSet to use the storage available on each of the running application pods to store the files locally
- O. Use a third-party tool to back up the EKS cluster for 1 year.

Answer: A

Explanation:

In the past, EBS can be attached only to one EC2 instance but not anymore but there are limitations like - it works only on io1/io2 instance types and many others as described here. <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volumes-multi.html> EFS has shareable storage. In terms of performance, Amazon EFS is optimized for workloads that require high levels of aggregate throughput and IOPS, whereas EBS is optimized for low-latency, random access I/O operations. Amazon EFS is designed to scale throughput and capacity automatically as your storage needs grow, while EBS volumes can be resized on demand.

NEW QUESTION 67

- (Exam Topic 2)

A company wants to run a custom network analysis software package to inspect traffic as traffic leaves and enters a VPC. The company has deployed the solution by using AWS CloudFormation on three Amazon EC2 instances in an Auto Scaling group. All network routing has been established to direct traffic to the EC2 instances.

Whenever the analysis software stops working, the Auto Scaling group replaces an instance. The network routes are not updated when the instance replacement occurs.

Which combination of steps will resolve this issue? (Select THREE.)

- A. Create alarms based on EC2 status check metrics that will cause the Auto Scaling group to replace the failed instance.
- B. Update the CloudFormation template to install the Amazon CloudWatch agent on the EC2 instances. Configure the CloudWatch agent to send process metrics for the application.
- C. Update the CloudFormation template to install AWS Systems Manager Agent on the EC2 instances. Configure Systems Manager Agent to send process metrics for the application.
- D. Create an alarm for the custom metric in Amazon CloudWatch for the failure scenario
- E. Configure the alarm to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.
- F. Create an AWS Lambda function that responds to the Amazon Simple Notification Service (Amazon SNS) message to take the instance out of service
- G. Update the network routes to point to the replacement instance.
- H. In the CloudFormation template, write a condition that updates the network routes when a replacement instance is launched.

Answer: BDE

NEW QUESTION 72

- (Exam Topic 2)

A company uses AWS Organizations for a multi-account setup in the AWS Cloud. The company's finance team has a data processing application that uses AWS Lambda and Amazon DynamoDB. The company's marketing team wants to access the data that is stored in the DynamoDB table. The DynamoDB table contains confidential data. The marketing team can have access to only specific attributes of data in the DynamoDB table. The finance team and the marketing team have separate AWS accounts.

What should a solutions architect do to provide the marketing team with the appropriate access to the DynamoDB table?

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

Answer: C

Explanation:

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

- Creating an SCP to grant the marketing team's AWS account access to the specific attributes of the DynamoDB table would not work because SCPs are policies that you can use with AWS Organizations to manage permissions in your organization's accounts. SCPs do not grant permissions; instead, they specify the maximum permissions that identities in an account can have². SCPs cannot be used to specify fine-grained access control for DynamoDB items and attributes.
- Creating an IAM role in the finance team's account by using IAM policy conditions for specific DynamoDB attributes and establishing trust with the marketing team's account would not work because IAM roles are identities that you can create in your account that have specific permissions. You can use an IAM role to delegate access to users, applications, or services that don't normally have access to your AWS resources³. However, creating an IAM role in the finance team's account would not restrict access to specific attributes of the DynamoDB table; it would only allow cross-account access. The company would still need a resource-based policy attached to the table to enforce fine-grained access control.
- Creating an IAM role in the finance team's account to access the DynamoDB table and using an IAM permissions boundary to limit the access to the specific attributes would not work because IAM permissions boundaries are policies that you use to delegate permissions management to other users. You can use permissions boundaries to limit the maximum permissions that an identity-based policy can grant to an IAM entity (user or role)⁴. Permissions boundaries cannot be used to specify fine-grained access control for DynamoDB items and attributes.

References:

- <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/specifying-conditions.html>
- https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html
- https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html
- https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_boundaries.html

NEW QUESTION 74

- (Exam Topic 2)

A company wants to containerize a multi-tier web application and move the application from an on-premises data center to AWS. The application includes web, application, and database tiers. The company needs to make the application fault tolerant and scalable. Some frequently accessed data must always be available across application servers. Frontend web servers need session persistence and must scale to meet increases in traffic.

Which solution will meet these requirements with the LEAST ongoing operational overhead?

- A. Run the application on Amazon Elastic Container Service (Amazon ECS) on AWS Fargate
- B. Use Amazon Elastic File System (Amazon EFS) for data that is frequently accessed between the web and application tier
- C. Store the frontend web server session data in Amazon Simple Queue Service (Amazon SQS).
- D. Run the application on Amazon Elastic Container Service (Amazon ECS) on Amazon EC2. Use Amazon ElastiCache for Redis to cache frontend web server session data
- E. Use Amazon Elastic Block Store (Amazon EBS) with Multi-Attach on EC2 instances that are distributed across multiple Availability Zones.
- F. Run the application on Amazon Elastic Kubernetes Service (Amazon EKS). Configure Amazon EKS to use managed node group
- G. Use ReplicaSets to run the web servers and application
- H. Create an Amazon Elastic File System (Amazon EFS) file system
- I. Mount the EFS file system across all EKS pods to store frontend web server session data.
- J. Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) Configure Amazon EKS to use managed node group
- K. Run the web servers and application as Kubernetes deployments in the EKS cluster
- L. Store the frontend web server session data in an Amazon DynamoDB table
- M. Create an Amazon Elastic File System (Amazon EFS) volume that all applications will mount at the time of deployment.

Answer: D

Explanation:

Deploying the application on Amazon EKS with managed node groups simplifies the operational overhead of managing the Kubernetes cluster. Running the web servers and application as Kubernetes deployments ensures that the desired number of pods are always running and can scale up or down as needed. Storing the frontend web server session data in an Amazon DynamoDB table provides a fast, scalable, and durable storage option that can be accessed across multiple Availability Zones. Creating an Amazon EFS volume that all applications will mount at the time of deployment allows the application to share data that is frequently accessed between the web and application tiers. References:

- <https://docs.aws.amazon.com/eks/latest/userguide/managed-node-groups.html>
- <https://docs.aws.amazon.com/eks/latest/userguide/deployments.html>
- <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html>
- <https://docs.aws.amazon.com/efs/latest/ug/mounting-fs.html>

NEW QUESTION 75

- (Exam Topic 2)

A company has five development teams that have each created five AWS accounts to develop and host applications. To track spending, the development teams log in to each account every month, record the current cost from the AWS Billing and Cost Management console, and provide the information to the company's finance team.

The company has strict compliance requirements and needs to ensure that resources are created only in AWS Regions in the United States. However, some resources have been created in other Regions.

A solutions architect needs to implement a solution that gives the finance team the ability to track and consolidate expenditures for all the accounts. The solution also must ensure that the company can create resources only in Regions in the United States.

Which combination of steps will meet these requirements in the MOST operationally efficient way? (Select THREE.)

- A. Create a new account to serve as a management account
- B. Create an Amazon S3 bucket for the finance team. Use AWS Cost and Usage Reports to create monthly reports and to store the data in the finance team's S3 bucket.
- C. Create a new account to serve as a management account
- D. Deploy an organization in AWS Organizations with all features enabled
- E. Invite all the existing accounts to the organization
- F. Ensure that each account accepts the invitation.
- G. Create an OU that includes all the development teams
- H. Create an SCP that allows the creation of resources only in Regions that are in the United States
- I. Apply the SCP to the OU.
- J. Create an OU that includes all the development teams
- K. Create an SCP that denies the creation of resources in Regions that are outside the United States
- L. Apply the SCP to the OU.
- M. Create an IAM role in the management account. Attach a policy that includes permissions to view the Billing and Cost Management console
- N. Allow the finance team users to assume the role
- O. Use AWS Cost Explorer and the Billing and Cost Management console to analyze costs.
- P. Create an IAM role in each AWS account
- Q. Attach a policy that includes permissions to view the Billing and Cost Management console
- R. Allow the finance team users to assume the role.

Answer: BCE

Explanation:

AWS Organizations is a service that enables you to consolidate multiple AWS accounts into an organization that you create and centrally manage. By creating a management account and inviting all the existing accounts to join the organization, the solutions architect can track and consolidate expenditures for all the accounts using AWS Cost Management tools such as AWS Cost Explorer and AWS Budgets. An organizational unit (OU) is a group of accounts within an organization that can be used to apply policies and simplify management. A service control policy (SCP) is a type of policy that you can use to manage permissions in your organization. By creating an OU that includes all the development teams and applying an SCP that allows the creation of resources only in Regions that are in the United States, the solutions architect can ensure that the company meets its compliance requirements and avoids unwanted charges from other Regions. An IAM role is an identity with permission policies that determine what the identity can and cannot do in AWS. By creating an IAM role in the management account and allowing the finance team users to assume it, the solutions architect can give them access to view the Billing and Cost Management console without sharing credentials or creating additional users. References:

- https://docs.aws.amazon.com/organizations/latest/userguide/orgs_introduction.html
- https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp.html
- https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html
- <https://docs.aws.amazon.com/aws-cost-management/latest/userguide/what-is-costmanagement.html>

NEW QUESTION 79

- (Exam Topic 2)

A company provides auction services for artwork and has users across North America and Europe. The company hosts its application in Amazon EC2 instances in the us-east-1 Region. Artists upload photos of their work as large-size, high-resolution image files from their mobile phones to a centralized Amazon S3 bucket created in the us-east-1 Region. The users in Europe are reporting slow performance for their image uploads.

How can a solutions architect improve the performance of the image upload process?

- A. Redeploy the application to use S3 multipart uploads.
- B. Create an Amazon CloudFront distribution and point to the application as a custom origin
- C. Configure the buckets to use S3 Transfer Acceleration.
- D. Create an Auto Scaling group for the EC2 instances and create a scaling policy.

Answer: C

Explanation:

Transfer acceleration. S3 Transfer Acceleration utilizes the Amazon CloudFront global network of edge locations to accelerate the transfer of data to and from S3 buckets. By enabling S3 Transfer Acceleration on the centralized S3 bucket, the users in Europe will experience faster uploads as their data will be routed through the closest CloudFront edge location.

NEW QUESTION 84

- (Exam Topic 2)

A company processes environment data. The company has a set up sensors to provide a continuous stream of data from different areas in a city. The data is available in JSON format.

The company wants to use an AWS solution to send the data to a database that does not require fixed schemas for storage. The data must be sent in real time.

Which solution will meet these requirements?

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

Answer: B

Explanation:

Amazon Kinesis Data Streams is a service that enables real-time data ingestion and processing. Amazon DynamoDB is a NoSQL database that does not require fixed schemas for storage. By using Kinesis Data Streams and DynamoDB, the company can send the JSON data to a database that can handle schemaless data in real time. References:

- > <https://docs.aws.amazon.com/streams/latest/dev/introduction.html>
- > <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html>

NEW QUESTION 86

- (Exam Topic 2)

A company needs to audit the security posture of a newly acquired AWS account. The company's data security team requires a notification only when an Amazon S3 bucket becomes publicly exposed. The company has already established an Amazon Simple Notification Service (Amazon SNS) topic that has the data security team's email address subscribed.

Which solution will meet these requirements?

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

Answer: B

Explanation:

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

NEW QUESTION 90

- (Exam Topic 1)

A company is running an event ticketing platform on AWS and wants to optimize the platform's cost-effectiveness. The platform is deployed on Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 and is backed by an Amazon RDS for MySQL DB instance. The company is developing new application features to run on Amazon EKS with AWS Fargate.

The platform experiences infrequent high peaks in demand. The surges in demand depend on event dates. Which solution will provide the MOST cost-effective setup for the platform?

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

Answer: B

Explanation:

They all mention using spot instances and EKS based on EC2. A spot instance is not appropriate for a production server and the company is developing new application designed for AWS Fargate, which means we must plan the future cost improvement including AWS Fargate.

<https://aws.amazon.com/savingsplans/compute-pricing/>

NEW QUESTION 95

- (Exam Topic 1)

A company is hosting an image-processing service on AWS in a VPC. The VPC extends across two Availability Zones. Each Availability Zone contains one public subnet and one private subnet.

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

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

Which solution will meet these requirements?

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

Answer: C

Explanation:

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

NEW QUESTION 99

- (Exam Topic 1)

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

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

What should the solutions architect recommend to meet these requirements?

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

Answer: A

Explanation:

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

NEW QUESTION 101

- (Exam Topic 1)

A company is developing a new serverless API by using Amazon API Gateway and AWS Lambda. The company integrated the Lambda functions with API Gateway to use several shared libraries and custom classes.

A solutions architect needs to simplify the deployment of the solution and optimize for code reuse. Which solution will meet these requirements?

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

Answer: B

Explanation:

Deploying the shared libraries and custom classes to a Docker image and uploading the image to Amazon Elastic Container Registry (Amazon ECR) and creating a Lambda layer that uses the Docker image as the source. Then, deploying the API's Lambda functions as Zip packages and configuring the packages to use the Lambda layer would meet the requirements for simplifying the deployment and optimizing for code reuse.

A Lambda layer is a distribution mechanism for libraries, custom runtimes, and other function dependencies. It allows you to manage your in-development function code separately from your dependencies, this way you can easily update your dependencies without having to update your entire function code.

By deploying the shared libraries and custom classes to a Docker image and uploading the image to Amazon Elastic Container Registry (ECR), it makes it easy to manage and version the dependencies. This way, the company can use the same version of the dependencies across different Lambda functions.

By creating a Lambda layer that uses the Docker image as the source, the company can configure the API's Lambda functions to use the layer, reducing the need

to include the dependencies in each function package, and making it easy to update the dependencies across all functions at once.

Reference:

AWS Lambda Layers documentation: <https://docs.aws.amazon.com/lambda/latest/dg/configuration-layers.html>

AWS Elastic Container Registry (ECR) documentation: <https://aws.amazon.com/ecr/> Building Lambda Layers with Docker documentation:

<https://aws.amazon.com/blogs/compute/building-lambda-layers-with-docker/>

NEW QUESTION 102

- (Exam Topic 1)

A company runs a content management application on a single Windows Amazon EC2 instance in a development environment. The application reads and writes static content to a 2 TB Amazon Elastic Block Store (Amazon EBS) volume that is attached to the instance as the root device. The company plans to deploy this application in production as a highly available and fault-tolerant solution that runs on at least three EC2 instances across multiple Availability Zones.

A solutions architect must design a solution that joins all the instances that run the application to an Active Directory domain. The solution also must implement Windows ACLs to control access to file contents. The application always must maintain exactly the same content on all running instances at any given point in time.

Which solution will meet these requirements with the LEAST management overhead?

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

Answer: C

Explanation:

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

NEW QUESTION 107

- (Exam Topic 1)

A large mobile gaming company has successfully migrated all of its on-premises infrastructure to the AWS Cloud. A solutions architect is reviewing the environment to ensure that it was built according to the design and that it is running in alignment with the Well-Architected Framework.

While reviewing previous monthly costs in Cost Explorer, the solutions architect notices that the creation and subsequent termination of several large instance types account for a high proportion of the costs. The solutions architect finds out that the company's developers are launching new Amazon EC2 instances as part of their testing and that the developers are not using the appropriate instance types.

The solutions architect must implement a control mechanism to limit the instance types that only the developers can launch.

Which solution will meet these requirements?

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

Answer: C

Explanation:

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

NEW QUESTION 111

- (Exam Topic 1)

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

What is the MOST cost-effective migration recommendation?

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

- I. Create a queue using Amazon M
- J. Configure the existing web server to publish to the new queue. When there are messages in the queue, invoke an AWS Lambda function to pull requests from the queue and process the file
- K. Store the processed files in Amazon EFS.
- L. Create a queue using Amazon SQS
- M. Configure the existing web server to publish to the new queue
- N. Use Amazon EC2 instances in an EC2 Auto Scaling group to pull requests from the queue and process the file
- O. Scale the EC2 instances based on the SQS queue length
- P. Store the processed files in an Amazon S3 bucket.

Answer: D

Explanation:

<https://aws.amazon.com/blogs/compute/operating-lambda-performance-optimization-part-1/>

NEW QUESTION 114

- (Exam Topic 1)

A company is refactoring its on-premises order-processing platform in the AWS Cloud. The platform includes a web front end that is hosted on a fleet of VMs RabbitMQ to connect the front end to the backend, and a Kubernetes cluster to run a containerized backend system to process the orders. The company does not want to make any major changes to the application

Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: A

Explanation:

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

NEW QUESTION 115

- (Exam Topic 1)

A company has 10 accounts that are part of an organization in AWS Organizations AWS Config is configured in each account All accounts belong to either the Prod OU or the NonProd OU

The company has set up an Amazon EventBridge rule in each AWS account to notify an Amazon Simple Notification Service (Amazon SNS) topic when an Amazon EC2 security group inbound rule is created with 0.0.0.0/0 as the source The company's security team is subscribed to the SNS topic

For all accounts in the NonProd OU the security team needs to remove the ability to create a security group inbound rule that includes 0.0.0.0/0 as the source

Which solution will meet this requirement with the LEAST operational overhead?

- A. Modify the EventBridge rule to invoke an AWS Lambda function to remove the security group inbound rule and to publish to the SNS topic Deploy the updated rule to the NonProd OU
- B. Add the vpc-sg-open-only-to-authorized-ports AWS Config managed rule to the NonProd OU
- C. Configure an SCP to allow the ec2:AuthorizeSecurityGroupIngress action when the value of the aws:SourceIp condition key is not 0.0.0.0/0 Apply the SCP to the NonProd OU
- D. Configure an SCP to deny the ec2:AuthorizeSecurityGroupIngress action when the value of the aws:SourceIp condition key is 0.0.0.0/0 Apply the SCP to the NonProd OU

Answer: D

Explanation:

This solution will meet the requirement with the least operational overhead because it directly denies the creation of the security group inbound rule with 0.0.0.0/0 as the source, which is the exact requirement. Additionally, it does not require any additional steps or resources such as invoking a Lambda function or adding a Config rule.

An SCP (Service Control Policy) is a policy that you can use to set fine-grained permissions for your AWS accounts within your organization. You can use SCPs to set permissions for the root user of an account and to delegate permissions to IAM users and roles in the accounts. You can use SCPs to set permissions that allow or deny access to specific services, actions, and resources.

To implement this solution, you would need to create an SCP that denies the ec2:AuthorizeSecurityGroupIngress action when the value of the aws:SourceIp condition key is 0.0.0.0/0. This SCP would then be applied to the NonProd OU. This would ensure that any security group inbound rule that includes 0.0.0.0/0 as the source will be denied, thus meeting the requirement.

Reference: https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp.html

https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_condition-keys.html

NEW QUESTION 120

- (Exam Topic 1)

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

Which steps should the solutions architect take to design an appropriate solution?

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

Availability Zone

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

Answer: C

Explanation:

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

NEW QUESTION 122

- (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 for cloud infrastructure to be allocated to the owning project. The team responsible for all of the AWS accounts has discovered that several Amazon EC2 instances are lacking the Project tag used for cost allocation.

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

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

Answer: ABE

Explanation:

<https://docs.aws.amazon.com/config/latest/developerguide/config-rule-multi-account-deployment.html>

<https://docs.aws.amazon.com/config/latest/developerguide/aggregate-data.html>

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

NEW QUESTION 126

- (Exam Topic 1)

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

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

Answer: AE

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2017/11/announcing-support-for-inter-region-vpc-peering/>

NEW QUESTION 127

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AWS-Certified-Solutions-Architect-Professional Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AWS-Certified-Solutions-Architect-Professional Product From:

<https://www.2passeasy.com/dumps/AWS-Certified-Solutions-Architect-Professional/>

Money Back Guarantee

AWS-Certified-Solutions-Architect-Professional Practice Exam Features:

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