



Amazon

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

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 needs to optimize the cost of backups for Amazon Elastic File System (Amazon EFS). A solutions architect has already configured a backup plan in AWS Backup for the EFS backups. The backup plan contains a rule with a lifecycle configuration to transition EFS backups to cold storage after 7 days and to keep the backups for an additional 90 days.

After 1 month, the company reviews its EFS storage costs and notices an increase in the EFS backup costs. The EFS backup cold storage produces almost double the cost of the EFS warm backup storage.

What should the solutions architect do to optimize the cost?

- A. Modify the backup rule's lifecycle configuration to move the EFS backups to cold storage after 1 day. Set the backup retention period to 30 days.
- B. Modify the backup rule's lifecycle configuration to move the EFS backups to cold storage after 8 days. Set the backup retention period to 30 days.
- C. Modify the backup rule's lifecycle configuration to move the EFS backups to cold storage after 1 day. Set the backup retention period to 90 days.
- D. Modify the backup rule's lifecycle configuration to move the EFS backups to cold storage after 8 days. Set the backup retention period to 98 days.

Answer: A

Explanation:

The cost of EFS backup cold storage is \$0.01 per GB-month, whereas the cost of EFS backup warm storage is \$0.05 per GB-month¹. Therefore, moving the backups to cold storage as soon as possible will reduce the storage cost. However, cold storage backups must be retained for a minimum of 90 days², otherwise they incur a pro-rated charge equal to the storage charge for the remaining days¹. Therefore, setting the backup retention period to 30 days will incur a penalty of 60 days of cold storage cost for each backup deleted. This penalty will still be lower than keeping the backups in warm storage for 7 days and then in cold storage for 83 days, which is the current configuration. Therefore, option A is the most cost-effective solution.

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 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 5

- (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 6

- (Exam Topic 2)

A company has IoT sensors that monitor traffic patterns throughout a large city. The company wants to read and collect data from the sensors and perform aggregations on the data.

A solutions architect designs a solution in which the IoT devices are streaming to Amazon Kinesis Data Streams. Several applications are reading from the stream. However, several consumers are experiencing throttling and are periodically encountering a RealProvisioned Throughput Exceeded error. Which actions should the solution architect take to resolve this issue? (Select THREE.)

- A. Reshard the stream to increase the number of shards in the stream.
- B. Use the Kinesis Producer Library (KPL). Adjust the polling frequency.

- C. Use consumers with the enhanced fan-out feature.
- D. Reshard the stream to reduce the number of shards in the stream.
- E. Use an error retry and exponential backoff mechanism in the consumer logic.
- F. Configure the stream to use dynamic partitioning.

Answer: ACE

Explanation:

<https://repost.aws/knowledge-center/kinesis-readprovisionedthroughputexceeded> Follow Data Streams best practices

To mitigate ReadProvisionedThroughputExceeded exceptions, apply these best practices:

- Reshard your stream to increase the number of shards in the stream.
- Use consumers with enhanced fan-out. For more information about enhanced fan-out, see Developing custom consumers with dedicated throughput (enhanced fan-out).
- Use an error retry and exponential backoff mechanism in the consumer logic if ReadProvisionedThroughputExceeded exceptions are encountered. For consumer applications that use an AWS SDK, the requests are retried by default.

NEW QUESTION 7

- (Exam Topic 2)

A solutions architect is redesigning a three-tier application that a company hosts on premises. The application provides personalized recommendations based on user profiles. The company already has an AWS account and has configured a VPC to host the application.

The frontend is a Java-based application that runs in on-premises VMs. The company hosts a personalization model on a physical application server and uses TensorFlow to implement the model. The personalization model uses artificial intelligence and machine learning (AI/ML). The company stores user information in a Microsoft SQL Server database. The web application calls the personalization model, which reads the user profiles from the database and provides recommendations.

The company wants to migrate the redesigned application to AWS.

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

- A. Use AWS Server Migration Service (AWS SMS) to migrate the on-premises physical application server and the web application VMs to AWS
- B. Use AWS Database Migration Service (AWS DMS) to migrate the SQL Server database to Amazon RDS for SQL Server.
- C. Export the personalization model
- D. Store the model artifacts in Amazon S3. Deploy the model to Amazon SageMaker and create an endpoint
- E. Host the Java application in AWS Elastic Beanstalk
- F. Use AWS Database Migration Service (AWS DMS) to migrate the SQL Server database to Amazon RDS for SQL Server.
- G. Use AWS Application Migration Service to migrate the on-premises personalization model and VMs to Amazon EC2 instances in Auto Scaling group
- H. Use AWS Database Migration Service (AWS DMS) to migrate the SQL Server database to an EC2 instance.
- I. Containerize the personalization model and the Java application
- J. Use Amazon Elastic Kubernetes Service (Amazon EKS) managed node groups to deploy the model and the application to Amazon EKS. Host the node groups in a VPC
- K. Use AWS Database Migration Service (AWS DMS) to migrate the SQL Server database to Amazon RDS for SQL Server.

Answer: B

Explanation:

Amazon SageMaker is a fully managed machine learning service that allows users to build, train, and deploy machine learning models quickly and easily¹. Users can export their existing TensorFlow models and store the model artifacts in Amazon S3, a highly scalable and durable object storage service². Users can then deploy the model to Amazon SageMaker and create an endpoint that can be invoked by the web application to provide recommendations³. This way, the solution can leverage the AI/ML capabilities of Amazon SageMaker without having to rewrite the personalization model.

AWS Elastic Beanstalk is a service that allows users to deploy and manage web applications without worrying about the infrastructure that runs those applications. Users can host their Java application in AWS Elastic Beanstalk and configure it to communicate with the Amazon SageMaker endpoint. This way, the solution can reduce the operational overhead of managing servers, load balancers, scaling, and application health monitoring.

AWS Database Migration Service (AWS DMS) is a service that helps users migrate databases to AWS quickly and securely. Users can use AWS DMS to migrate their SQL Server database to Amazon RDS for SQL Server, a fully managed relational database service that offers high availability, scalability, security, and compatibility. This way, the solution can reduce the operational overhead of managing database servers, backups, patches, and upgrades.

Option A is incorrect because using AWS Server Migration Service (AWS SMS) to migrate the on-premises physical application server and the web application VMs to AWS is not cost-effective or scalable. AWS SMS is a service that helps users migrate on-premises workloads to AWS. However, for this use case, migrating the physical application server and the web application VMs to AWS will not take advantage of the AI/ML capabilities of Amazon SageMaker or the managed services of AWS Elastic Beanstalk and Amazon RDS.

Option C is incorrect because using AWS Application Migration Service to migrate the on-premises personalization model and VMs to Amazon EC2 instances in Auto Scaling groups is not cost-effective or scalable. AWS Application Migration Service is a service that helps users migrate applications from on-premises or other clouds to AWS without making any changes to their applications. However, for this use case, migrating the personalization model and VMs to EC2 instances will not take advantage of the AI/ML capabilities of Amazon SageMaker or the managed services of AWS Elastic Beanstalk and Amazon RDS.

Option D is incorrect because containerizing the personalization model and the Java application and using Amazon Elastic Kubernetes Service (Amazon EKS) managed node groups to deploy them to Amazon EKS is not necessary or cost-effective. Amazon EKS is a service that allows users to run Kubernetes on AWS without needing to install, operate, and maintain their own Kubernetes control plane or nodes. However, for this use case, containerizing and deploying the personalization model and the Java application will not take advantage of the AI/ML capabilities of Amazon SageMaker or the managed services of AWS Elastic Beanstalk. Moreover, using S3 Glacier Deep Archive as a storage class for images will incur a high retrieval fee and latency for accessing them.

NEW QUESTION 8

- (Exam Topic 2)

A company wants to use AWS for disaster recovery for an on-premises application. The company has hundreds of Windows-based servers that run the application. All the servers mount a common share.

The company has an RTO of 15 minutes and an RPO of 5 minutes. The solution must support native failover and fallback capabilities.

Which solution will meet these requirements MOST cost-effectively?

- A. Create an AWS Storage Gateway File Gateway
- B. Schedule daily Windows server backup
- C. Save the data to Amazon S3. During a disaster, recover the on-premises servers from the backup
- D. During failback
- E. run the on-premises servers on Amazon EC2 instances.
- F. Create a set of AWS CloudFormation templates to create infrastructure
- G. Replicate all data to Amazon Elastic File System (Amazon EFS) by using AWS DataSync

- H. During a disaster, use AWS CodePipeline to deploy the templates to restore the on-premises server
- I. Fail back the data by using DataSync.
- J. Create an AWS Cloud Development Kit (AWS CDK) pipeline to stand up a multi-site active-active environment on AW
- K. Replicate data into Amazon S3 by using the s3 sync comman
- L. During a disaster, swap DNS endpoints to point to AW
- M. Fail back the data by using the s3 sync command.
- N. Use AWS Elastic Disaster Recovery to replicate the on-premises server
- O. Replicate data to an Amazon FSx for Windows File Server file system by using AWS DataSyn
- P. Mount the file system to AWS server
- Q. During a disaster, fail over the on-premises servers to AW
- R. Fail back to new or existing servers by using Elastic Disaster Recovery.

Answer: D

NEW QUESTION 9

- (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 instanc
- B. Store the connection credentials as a secret in AWS Secrets Manager.
- C. Deploy an Amazon RDS Proxy layer in front of the DB instanc
- D. Store the connection credentials in AWS Systems Manager Parameter Store.
- E. Create an Aurora Replic
- F. Store the connection credentials as a secret in AWS Secrets Manager.
- G. Create an Aurora Replic
- 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 10

- (Exam Topic 2)

A company is using AWS CloudFormation to deploy its infrastructure. The company is concerned that, if a production CloudFormation stack is deleted, important data stored in Amazon RDS databases or Amazon EBS volumes might also be deleted.

How can the company prevent users from accidentally deleting data in this way?

- A. Modify the CloudFormation templates to add a DeletionPolicy attribute to RDS and EBS resources.
- B. Configure a stack policy that disallows the deletion of RDS and EBS resources.
- C. Modify IAM policies to deny deleting RDS and EBS resources that are tagged with an "awscloudformation: stack-name" tag.
- D. Use AWS Config rules to prevent deleting RDS and EBS resources.

Answer: A

Explanation:

With the DeletionPolicy attribute you can preserve or (in some cases) backup a resource when its stack is deleted. You specify a DeletionPolicy attribute for each resource that you want to control. If a resource has no DeletionPolicy attribute, AWS CloudFormation deletes the resource by default. To keep a resource when its stack is deleted, specify Retain for that resource. You can use retain for any resource. For example, you can retain a nested stack, Amazon S3 bucket, or EC2 instance so that you can continue to use or modify those resources after you delete their stacks.

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-deletionpolicy.html>

NEW QUESTION 10

- (Exam Topic 2)

A company has an application in the AWS Cloud. The application runs on a fleet of 20 Amazon EC2 instances. The EC2 instances are persistent and store data on multiple attached Amazon Elastic Block Store (Amazon EBS) volumes.

The company must maintain backups in a separate AWS Region. The company must be able to recover the EC2 instances and their configuration within 1 business day, with loss of no more than 1 day's worth of data. The company has limited staff and needs a backup solution that optimizes operational efficiency and cost. The company already has created an AWS CloudFormation template that can deploy the required network configuration in a secondary Region.

Which solution will meet these requirements?

- A. Create a second CloudFormation template that can recreate the EC2 instances in the secondary Region. Run daily multivolume snapshots by using AWS Systems Manager Automation runbook
- B. Copy the snapshots to the secondary Regio
- C. In the event of a failure, launch the CloudFormation templates, restore the EBS volumes from snapshots, and transfer usage to the secondary Region.
- D. Use Amazon Data Lifecycle Manager (Amazon DLM) to create daily multivolume snapshots of the EBS volume
- E. In the event of a failure, launch the CloudFormation template and use Amazon DLM to restore the EBS volumes and transfer usage to the secondary Region.
- F. Use AWS Backup to create a scheduled daily backup plan for the EC2 instance
- G. Configure the backup task to copy the backups to a vault in the secondary regio
- H. In the event of a failure, launch the CloudFormation template, restore the instance volumes and configurations from the backup vault, and transfer usage to the secondary Region.
- I. Deploy EC2 instances of the same size and configuration to the secondary Regio
- J. Configure AWS DataSync daily to copy data from the primary Region to the secondary Regio
- K. In the event of a failure, launch the CloudFormation template and transfer usage to the secondary Region.

Answer: C

Explanation:

Using AWS Backup to create a scheduled daily backup plan for the EC2 instances will enable taking snapshots of the EC2 instances and their attached EBS volumes1. Configuring the backup task to copy the backups to a vault in the secondary Region will enable maintaining backups in a separate Region1. In the event of a failure, launching the CloudFormation template will enable deploying the network configuration in the secondary Region2. Restoring the instance volumes and configurations from the backup vault will enable recovering the EC2 instances and their data1. Transferring usage to the secondary Region will enable resuming operations2.

NEW QUESTION 14

- (Exam Topic 2)

A solutions architect needs to define a reference architecture for a solution for three-tier applications with web, application, and NoSQL data layers. The reference architecture must meet the following requirements:

- High availability within an AWS Region
- Able to fail over in 1 minute to another AWS Region for disaster recovery
- Provide the most efficient solution while minimizing the impact on the user experience Which combination of steps will meet these requirements? (Select THREE.)

- A. Use an Amazon Route 53 weighted routing policy set to 100/0 across the two selected Region
- B. Set Time to Live (TTL) to 1 hour.
- C. Use an Amazon Route 53 failover routing policy for failover from the primary Region to the disaster recovery Region
- D. Set Time to Live (TTL) to 30 seconds.
- E. Use a global table within Amazon DynamoDB so data can be accessed in the two selected Regions.
- F. Back up data from an Amazon DynamoDB table in the primary Region every 60 minutes and then write the data to Amazon S3. Use S3 Cross-Region replication to copy the data from the primary Region to the disaster recovery Region
- G. Have a script import the data into DynamoDB in a disaster recovery scenario.
- H. Implement a hot standby model using Auto Scaling groups for the web and application layers across multiple Availability Zones in the Region
- I. Use zonal Reserved Instances for the minimum number of servers and On-Demand Instances for any additional resources.
- J. Use Auto Scaling groups for the web and application layers across multiple Availability Zones in the Region
- K. Use Spot Instances for the required resources.

Answer: BCE

Explanation:

The requirements can be achieved by using an Amazon DynamoDB database with a global table. DynamoDB is a NoSQL database so it fits the requirements. A global table also allows both reads and writes to occur in both Regions. For the web and application tiers Auto Scaling groups should be configured. Due to the 1-minute RTO these must be configured in an active/passive state. The best pricing model to lower price but ensure resources are available when needed is to use a combination of zonal reserved instances and on-demand instances. To failover between the Regions, a Route 53 failover routing policy can be configured with a TTL configured on the record of 30 seconds. This will mean clients must resolve against Route 53 every 30 seconds to get the latest record. In a failover scenario the clients would be redirected to the secondary site if the primary site is unhealthy.

NEW QUESTION 15

- (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 minute1s. AWS DRS continuously replicates data from the source servers to a staging area subnet in the DR Region, where it uses low-cost

storage and minimal compute resources to maintain ongoing replication. In the event of a disaster, AWS DRS automatically converts the servers to boot and run natively on AWS and launches recovery instances on AWS within minutes². By using AWS DRS, the company can save costs by removing idle recovery site resources and paying for the full disaster recovery site only when needed. By creating a cross-Region read replica for the DB instance, the company can have a standby copy of its primary database in a different AWS Region³. By using infrastructure as code (IaC), the company can provision the new infrastructure in the DR Region in an automated and consistent way⁴. By using an Amazon Route 53 failover routing policy, the company can route traffic to a resource that is healthy or to another resource when the first resource becomes unavailable.

The other options are not correct because:

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

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

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

References:

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

NEW QUESTION 18

- (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 22

- (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 23

- (Exam Topic 2)

A company runs an IoT application in the AWS Cloud. The company has millions of sensors that collect data from houses in the United States. The sensors use the MQTT protocol to connect and send data to a custom MQTT broker. The MQTT broker stores the data on a single Amazon EC2 instance. The sensors connect to the broker through the domain named `iot.example.com`. The company uses Amazon Route 53 as its DNS service. The company stores the data in Amazon DynamoDB.

On several occasions, the amount of data has overloaded the MQTT broker and has resulted in lost sensor data. The company must improve the reliability of the solution.

Which solution will meet these requirements?

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

Answer: A

Explanation:

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

Reference: <https://aws.amazon.com/elasticloadbalancing/applicationloadbalancer/> <https://aws.amazon.com/autoscaling/> <https://aws.amazon.com/route53/>

NEW QUESTION 28

- (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 29

- (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 32

- (Exam Topic 2)

A company needs to migrate its customer transactions database from on premises to AWS. The database resides on an Oracle DB instance that runs on a Linux server. According to a new security requirement, the company must rotate the database password each year.

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

- A. Convert the database to Amazon DynamoDB by using the AWS Schema Conversion Tool (AWS SCT). Store the password in AWS Systems Manager Parameter Store
- B. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function for yearly password rotation.
- C. Migrate the database to Amazon RDS for Oracle
- D. Store the password in AWS Secrets Manager
- E. Turn on automatic rotation
- F. Configure a yearly rotation schedule.
- G. Migrate the database to an Amazon EC2 instance
- H. Use AWS Systems Manager Parameter Store to keep and rotate the connection string by using an AWS Lambda function on a yearly schedule
- I. Migrate the database to Amazon Neptune by using the AWS Schema Conversion Tool (AWS SCT). Create an Amazon CloudWatch alarm to invoke an AWS Lambda function for yearly password rotation.

Answer: B

NEW QUESTION 34

- (Exam Topic 2)

A company is implementing a serverless architecture by using AWS Lambda functions that need to access a Microsoft SQL Server DB instance on Amazon RDS. The company has separate environments for development and production, including a clone of the database system.

The company's developers are allowed to access the credentials for the development database. However, the credentials for the production database must be encrypted with a key that only members of the IT security team's IAM user group can access. This key must be rotated on a regular basis.

What should a solutions architect do in the production environment to meet these requirements?

- A. Store the database credentials in AWS Systems Manager Parameter Store by using a SecureString parameter that is encrypted by an AWS Key Management Service (AWS KMS) customer managed key
- B. Attach a role to each Lambda function to provide access to the SecureString parameter
- C. Restrict access to the SecureString parameter and the customer managed key so that only the IT security team can access the parameter and the key.
- D. Encrypt the database credentials by using the AWS Key Management Service (AWS KMS) default Lambda key
- E. Store the credentials in the environment variables of each Lambda function
- F. Load the credentials from the environment variables in the Lambda code
- G. Restrict access to the KMS key so that only the IT security team can access the key.
- H. Store the database credentials in the environment variables of each Lambda function
- I. Encrypt the environment variables by using an AWS Key Management Service (AWS KMS) customer managed key
- J. Restrict access to the customer managed key so that only the IT security team can access the key.
- K. Store the database credentials in AWS Secrets Manager as a secret that is associated with an AWS Key Management Service (AWS KMS) customer managed key
- L. Attach a role to each Lambda function to provide access to the secret
- M. Restrict access to the secret and the customer managed key so that only the IT security team can access the secret and the key.

Answer: D

Explanation:

Storing the database credentials in AWS Secrets Manager as a secret that is associated with an AWS Key Management Service (AWS KMS) customer managed key will enable encrypting and managing the credentials securely. 1. AWS Secrets Manager helps you to securely encrypt, store, and retrieve credentials for your databases and other services. 2. Attaching a role to each Lambda function to provide access to the secret will enable retrieving the credentials programmatically. 1. Restricting access to the secret and the customer managed key so that only members of the IT security team's IAM user group can access them will enable meeting the security requirements. 1.

NEW QUESTION 38

- (Exam Topic 2)

A company has a few AWS accounts for development and wants to move its production application to AWS. The company needs to enforce Amazon Elastic Block Store (Amazon EBS) encryption at rest on current production accounts and future production accounts only. The company needs a solution that includes built-in blueprints and guardrails.

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

- A. Use AWS CloudFormation StackSets to deploy AWS Config rules on production accounts.
- B. Create a new AWS Control Tower landing zone in an existing developer account
- C. Create OUs for account
- D. Add production and development accounts to production and development OUs, respectively.
- E. Create a new AWS Control Tower landing zone in the company's management account
- F. Add production and development accounts to production and development OU
- G. respectively.
- H. Invite existing accounts to join the organization in AWS Organization
- I. Create SCPs to ensure compliance.
- J. Create a guardrail from the management account to detect EBS encryption.
- K. Create a guardrail for the production OU to detect EBS encryption.

Answer: CDF

Explanation:

<https://docs.aws.amazon.com/controltower/latest/userguide/controls.html> <https://docs.aws.amazon.com/controltower/latest/userguide/strongly-recommended-controls.html#ebs-enable-en> AWS is now transitioning the previous term 'guardrail' new term 'control'.

NEW QUESTION 40

- (Exam Topic 2)

A software-as-a-service (SaaS) provider exposes APIs through an Application Load Balancer (ALB). The ALB connects to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster that is deployed in the us-east-1 Region. The exposed APIs contain usage of a few non-standard REST methods: LINK, UNLINK, LOCK, and UNLOCK.

Users outside the United States are reporting long and inconsistent response times for these APIs. A solutions architect needs to resolve this problem with a solution that minimizes operational overhead.

Which solution meets these requirements?

- A. Add an Amazon CloudFront distributio
- B. Configure the ALB as the origin.
- C. Add an Amazon API Gateway edge-optimized API endpoint to expose the API
- D. Configure the ALB as the target.
- E. Add an accelerator in AWS Global Accelerato
- F. Configure the ALB as the origin.
- G. Deploy the APIs to two additional AWS Regions: eu-west-1 and ap-southeast-2. Add latency-based routing records in Amazon Route 53.

Answer: C

Explanation:

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

NEW QUESTION 43

- (Exam Topic 2)

A company is migrating a legacy application from an on-premises data center to AWS. The application uses MongoDB as a key-value database. According to the company's technical guidelines, all Amazon EC2 instances must be hosted in a private subnet without an internet connection. In addition, all connectivity between applications and databases must be encrypted. The database must be able to scale based on demand.

Which solution will meet these requirements?

- A. Create new Amazon DocumentDB (with MongoDB compatibility) tables for the application with Provisioned IOPS volume
- B. Use the instance endpoint to connect to Amazon DocumentDB.
- C. Create new Amazon DynamoDB tables for the application with on-demand capacit
- D. Use a gateway VPC endpoint for DynamoDB to connect to the DynamoDB tables
- E. Create new Amazon DynamoDB tables for the application with on-demand capacit
- F. Use an interface VPC endpoint for DynamoDB to connect to the DynamoDB tables.
- G. Create new Amazon DocumentDB (with MongoDB compatibility) tables for the application with Provisioned IOPS volumes. Use the cluster endpoint to connect to Amazon DocumentDB

Answer: A

Explanation:

A is the correct answer because it uses Amazon DocumentDB (with MongoDB compatibility) as a key-value database that can scale based on demand and supports encryption in transit and at rest. Amazon DocumentDB is a fully managed document database service that is designed to be compatible with the MongoDB API. It is a NoSQL database that is optimized for storing, indexing, and querying JSON data. Amazon DocumentDB supports encryption in transit using TLS and encryption at rest using AWS Key Management Service (AWS KMS). Amazon DocumentDB also supports provisioned IOPS volumes that can scale up to 64 TiB of storage and 256,000 IOPS per cluster. To connect to Amazon DocumentDB, you can use the instance endpoint, which connects to a specific instance in the cluster, or the cluster endpoint, which connects to the primary instance or one of the replicas in the cluster. Using the cluster endpoint is recommended for high availability and load balancing purposes. References:

- > <https://docs.aws.amazon.com/documentdb/latest/developerguide/what-is.html>
- > <https://docs.aws.amazon.com/documentdb/latest/developerguide/security.encryption.html>
- > <https://docs.aws.amazon.com/documentdb/latest/developerguide/limits.html>
- > <https://docs.aws.amazon.com/documentdb/latest/developerguide/connecting.html>

NEW QUESTION 44

- (Exam Topic 2)

A solutions architect needs to review the design of an Amazon EMR cluster that is using the EMR File System (EMRFS). The cluster performs tasks that are critical to business needs. The cluster is running Amazon EC2 On-Demand Instances at all times for all task, primary, and core nodes. The EMR tasks run each morning, starting at 1 :00 AM. and take 6 hours to finish running. The amount of time to complete the processing is not a priority because the data is not referenced until late in the day.

The solutions architect must review the architecture and suggest a solution to minimize the compute costs. Which solution should the solutions architect recommend to meet these requirements?

- A. Launch all task, primary, and core nodes on Spot Instances in an instance fle
- B. Terminate the cluster, including all instances, when the processing is completed.
- C. Launch the primary and core nodes on On-Demand Instance
- D. Launch the task nodes on Spot Instances in an instance fle
- E. Terminate the cluster, including all instances, when the processing is complete
- F. Purchase Compute Savings Plans to cover the On-Demand Instance usage.
- G. Continue to launch all nodes on On-Demand Instance
- H. Terminate the cluster, including all instances, when the processing is complete
- I. Purchase Compute Savings Plans to cover the On-Demand Instance usage
- J. Launch the primary and core nodes on On-Demand Instance
- K. Launch the task nodes on Spot Instances in an instance fle
- L. Terminate only the task node instances when the processing is complete
- M. Purchase Compute Savings Plans to cover the On-Demand Instance usage.

Answer: A

Explanation:

Amazon EC2 Spot Instances offer spare compute capacity at steep discounts compared to On-Demand prices. Spot Instances can be interrupted by EC2 with two minutes of notification when EC2 needs the capacity back. Amazon EMR can handle Spot interruptions gracefully by decommissioning the nodes and redistributing the tasks to other nodes. By launching all nodes on Spot Instances in an instance fleet, the solutions architect can minimize the compute costs of the EMR cluster. An instance fleet is a collection of EC2 instances with different types and sizes that EMR automatically provisions to meet a defined target capacity. By terminating the cluster when the processing is completed, the solutions architect can avoid paying for idle resources. References:

- > <https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-managed-scaling.html>
- > <https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-instance-fleet.html>
- > <https://aws.amazon.com/blogs/big-data/optimizing-amazon-emr-for-resilience-and-cost-with-capacity-opt>

NEW QUESTION 47

- (Exam Topic 2)

A global manufacturing company plans to migrate the majority of its applications to AWS. However, the company is concerned about applications that need to remain within a specific country or in the company's central on-premises data center because of data regulatory requirements or requirements for latency of single-digit milliseconds. The company also is concerned about the applications that it hosts in some of its factory sites, where limited network infrastructure exists. The company wants a consistent developer experience so that its developers can build applications once and deploy on premises, in the cloud, or in a hybrid architecture.

The developers must be able to use the same tools, APIs, and services that are familiar to them. Which solution will provide a consistent hybrid experience to meet these requirements?

- A. Migrate all applications to the closest AWS Region that is compliant
- B. Set up an AWS Direct Connect connection between the central on-premises data center and AWS
- C. Deploy a Direct Connect gateway.
- D. Use AWS Snowball Edge Storage Optimized devices for the applications that have data regulatory requirements or requirements for latency of single-digit millisecond
- E. Retain the devices on premise
- F. Deploy AWS Wavelength to host the workloads in the factory sites.
- G. Install AWS Outposts for the applications that have data regulatory requirements or requirements for latency of single-digit millisecond
- H. Use AWS Snowball Edge Compute Optimized devices to host the workloads in the factory sites.
- I. Migrate the applications that have data regulatory requirements or requirements for latency of single-digit milliseconds to an AWS Local Zone
- J. Deploy AWS Wavelength to host the workloads in the factory sites.

Answer: C

Explanation:

Installing AWS Outposts for the applications that have data regulatory requirements or requirements for latency of single-digit milliseconds will provide a fully managed service that extends AWS infrastructure, services, APIs, and tools to customer premises¹. AWS Outposts allows customers to run some AWS services locally and connect to a broad range of services available in the local AWS Region¹. Using AWS Snowball Edge Compute Optimized devices to host the workloads in the factory sites will provide local compute and storage resources for locations with limited network infrastructure². AWS Snowball Edge devices can run Amazon EC2 instances and AWS Lambda functions locally and sync data with AWS when network connectivity is available².

NEW QUESTION 52

- (Exam Topic 2)

A company is running an application on Amazon EC2 instances in the AWS Cloud. The application is using a MongoDB database with a replica set as its data tier. The MongoDB database is installed on systems in the company's on-premises data center and is accessible through an AWS Direct Connect connection to the data center environment.

A solutions architect must migrate the on-premises MongoDB database to Amazon DocumentDB (with MongoDB compatibility).

Which strategy should the solutions architect choose to perform this migration?

- A. Create a fleet of EC2 instances
- B. Install MongoDB Community Edition on the EC2 instances, and create a database
- C. Configure continuous synchronous replication with the database that is running in the on-premises data center.
- D. Create an AWS Database Migration Service (AWS DMS) replication instance
- E. Create a source endpoint for the on-premises MongoDB database by using change data capture (CDC). Create a target endpoint for the Amazon DocumentDB database
- F. Create and run a DMS migration task.
- G. Create a data migration pipeline by using AWS Data Pipeline
- H. Define data nodes for the on-premises MongoDB database and the Amazon DocumentDB database
- I. Create a scheduled task to run the data pipeline.
- J. Create a source endpoint for the on-premises MongoDB database by using AWS Glue crawlers. Configure continuous asynchronous replication between the MongoDB database and the Amazon DocumentDB database.

Answer: B

Explanation:

<https://aws.amazon.com/getting-started/hands-on/move-to-managed/migrate-mongodb-to-documentdb/>

NEW QUESTION 55

- (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[user_id] 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 58

- (Exam Topic 2)

A company is creating a REST API to share information with six of its partners based in the United States. The company has created an Amazon API Gateway Regional endpoint. Each of the six partners will access the API once per day to post daily sales figures.

After initial deployment, the company observes 1,000 requests per second originating from 500 different IP addresses around the world. The company believes this traffic is originating from a botnet and wants to secure its API while minimizing cost.

Which approach should the company take to secure its API?

- A. Create an Amazon CloudFront distribution with the API as the origin
- B. Create an AWS WAF web ACL with a rule to block clients that submit more than five requests per day
- C. Associate the web ACL with the CloudFront distribution
- D. Configure CloudFront with an origin access identity (OAI) and associate it with the distribution
- E. Configure API Gateway to ensure only the OAI can run the POST method.
- F. Create an Amazon CloudFront distribution with the API as the origin
- G. Create an AWS WAF web ACL with a rule to block clients that submit more than five requests per day
- H. Associate the web ACL with the CloudFront distribution
- I. Add a custom header to the CloudFront distribution populated with an API key
- J. Configure the API to require an API key on the POST method.
- K. Create an AWS WAF web ACL with a rule to allow access to the IP addresses used by the six partners. Associate the web ACL with the API
- L. Create a resource policy with a request limit and associate it with the API
- M. Configure the API to require an API key on the POST method.
- N. Create an AWS WAF web ACL with a rule to allow access to the IP addresses used by the six partners. Associate the web ACL with the API
- O. Create a usage plan with a request limit and associate it with the API
- P. Create an API key and add it to the usage plan.

Answer: D

Explanation:

"A usage plan specifies who can access one or more deployed API stages and methods—and also how much and how fast they can access them. The plan uses API keys to identify API clients and meters access to the associated API stages for each key. It also lets you configure throttling limits and quota limits that are enforced on individual client API keys."

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-usage-plans.html>

A rate-based rule tracks the rate of requests for each originating IP address, and triggers the rule action on IPs with rates that go over a limit. You set the limit as the number of requests per 5-minute time span..... The following caveats apply to AWS WAF rate-based rules: The minimum rate that you can set is 100. AWS WAF checks the rate of requests every 30 seconds, and counts requests for the prior five minutes each time. Because of this, it's possible for an IP address to send requests at too high a rate for 30 seconds before AWS WAF detects and blocks it. AWS WAF can block up to 10,000 IP addresses. If more than 10,000 IP addresses send high rates of requests at the same time, AWS WAF will only block 10,000 of them. " <https://docs.aws.amazon.com/waf/latest/developerguide/waf-rule-statement-type-rate-based.html>

NEW QUESTION 63

- (Exam Topic 2)

A solutions architect is reviewing a company's process for taking snapshots of Amazon RDS DB instances. The company takes automatic snapshots every day and retains the snapshots for 7 days.

The solutions architect needs to recommend a solution that takes snapshots every 6 hours and retains the snapshots for 30 days. The company uses AWS Organizations to manage all of its AWS accounts. The company needs a consolidated view of the health of the RDS snapshots.

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

- A. Turn on the cross-account management feature in AWS Backup
- B. Create a backup plan that specifies the frequency and retention requirement
- C. Add a tag to the DB instance
- D. Apply the backup plan by using tag
- E. Use AWS Backup to monitor the status of the backups.
- F. Turn on the cross-account management feature in Amazon RDS
- G. Create a snapshot global policy that specifies the frequency and retention requirement
- H. Use the RDS console in the management account to monitor the status of the backups.
- I. Turn on the cross-account management feature in AWS CloudFormation
- J. From the management account, deploy a CloudFormation stack set that contains a backup plan from AWS Backup that specifies the frequency and retention requirement
- K. Create an AWS Lambda function in the management account to monitor the status of the backup
- L. Create an Amazon EventBridge rule in each account to run the Lambda function on a schedule.
- M. Configure AWS Backup in each account
- N. Create an Amazon Data Lifecycle Manager lifecycle policy that specifies the frequency and retention requirement
- O. Specify the DB instances as the target resource
- P. Use the Amazon Data Lifecycle Manager console in each member account to monitor the status of the backups.

Answer: A

Explanation:

Turning on the cross-account management feature in AWS Backup will enable managing and monitoring backups across multiple AWS accounts that belong to the same organization in AWS Organizations¹. Creating a backup plan that specifies the frequency and retention requirements will enable taking snapshots every 6 hours and retaining them for 30 days². Adding a tag to the DB instances will enable applying the backup plan by using tags². Using AWS Backup to monitor the status of the backups will enable having a consolidated view of the health of the RDS snapshots¹.

NEW QUESTION 64

- (Exam Topic 2)

A company has multiple business units that each have separate accounts on AWS. Each business unit manages its own network with several VPCs that have CIDR ranges that overlap. The company's marketing team has created a new internal application and wants to make the application accessible to all the other business units. The solution must use private IP addresses only.

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

- A. Instruct each business unit to add a unique secondary CIDR range to the business unit's VPC
- B. Peer the VPCs and use a private NAT gateway in the secondary range to route traffic to the marketing team.
- C. Create an Amazon EC2 instance to serve as a virtual appliance in the marketing account's VPC
- D. Create an AWS Site-to-Site VPN connection between the marketing team and each business unit's VPC
- E. Perform NAT where necessary.
- F. Create an AWS PrivateLink endpoint service to share the marketing application
- G. Grant permission to specific AWS accounts to connect to the service
- H. Create interface VPC endpoints in other accounts to access the application by using private IP addresses.
- I. Create a Network Load Balancer (NLB) in front of the marketing application in a private subnet
- J. Create an API Gateway AP
- K. Use the Amazon API Gateway private integration to connect the API to the NLB
- L. Activate IAM authorization for the AP
- M. Grant access to the accounts of the other business units.

Answer: C

Explanation:

With AWS PrivateLink, the marketing team can create an endpoint service to share their internal application with other accounts securely using private IP addresses. They can grant permission to specific AWS accounts to connect to the service and create interface VPC endpoints in the other accounts to access the application by using private IP addresses. This option does not require any changes to the network of the other business units, and it does not require peering or NATing. This solution is both scalable and secure.

<https://aws.amazon.com/blogs/networking-and-content-delivery/connecting-networks-with-overlapping-ip-range>

NEW QUESTION 67

- (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-repository>

NEW QUESTION 70

- (Exam Topic 2)

A company has an on-premises Microsoft SQL Server database that writes a nightly 200 GB export to a local drive. The company wants to move the backups to more robust cloud storage on Amazon S3. The company has set up a 10 Gbps AWS Direct Connect connection between the on-premises data center and AWS. Which solution meets these requirements MOST cost-effectively?

- A. Create a new S3 bucket
- B. Deploy an AWS Storage Gateway file gateway within the VPC that is connected to the Direct Connect connection
- C. Create a new SMB file share
- D. Write nightly database exports to the new SMB file share.
- E. Create an Amazon FSx for Windows File Server Single-AZ file system within the VPC that is connected to the Direct Connect connection
- F. Create a new SMB file share
- G. Write nightly database exports to an SMB file share on the Amazon FSx file system
- H. Enable nightly backups.
- I. Create an Amazon FSx for Windows File Server Multi-AZ file system within the VPC that is connected to the Direct Connect connection
- J. Create a new SMB file share
- K. Write nightly database exports to an SMB file share on the Amazon FSx file system
- L. Enable nightly backups.
- M. Create a new S3 bucket
- N. Deploy an AWS Storage Gateway volume gateway within the VPC that is connected to the Direct Connect connection
- O. Create a new SMB file share

P. Write nightly database exports to the new SMB file share on the volume gateway, and automate copies of this data to an S3 bucket.

Answer: A

Explanation:

<https://docs.aws.amazon.com/filegateway/latest/files3/CreatingAnSMBFileShare.html>

NEW QUESTION 74

- (Exam Topic 2)

A company is running a containerized application in the AWS Cloud. The application is running by using Amazon Elastic Container Service (Amazon ECS) on a set of Amazon EC2 instances. The EC2 instances run in an Auto Scaling group.

The company uses Amazon Elastic Container Registry (Amazon ECR) to store its container images. When a new image version is uploaded, the new image version receives a unique tag.

The company needs a solution that inspects new image versions for common vulnerabilities and exposures. The solution must automatically delete new image tags that have Critical or High severity findings. The solution also must notify the development team when such a deletion occurs.

Which solution meets these requirements?

- A. Configure scan on push on the repository Use Amazon EventBridge to invoke an AWS Step Functions state machine when a scan is complete for images that have Critical or High severity finding
- B. Use the Step Functions state machine to delete the image tag for those images and to notify the development team through Amazon Simple Notification Service (Amazon SNS).
- C. Configure scan on push on the repository Configure scan results to be pushed to an Amazon Simple Queue Service (Amazon SQS) queue
- D. Invoke an AWS Lambda function when a new message is added to the SQS queue
- E. Use the Lambda function to delete the image tag for images that have Critical or High severity finding
- F. Notify the development team by using Amazon Simple Email Service (Amazon SES).
- G. Schedule an AWS Lambda function to start a manual image scan every hour
- H. Configure Amazon EventBridge to invoke another Lambda function when a scan is complete
- I. Use the second Lambda function to delete the image tag for images that have Critical or High severity finding
- J. Notify the development team by using Amazon Simple Notification Service (Amazon SNS).
- K. Configure periodic image scan on the repository
- L. Configure scan results to be added to an Amazon Simple Queue Service (Amazon SQS) queue
- M. Invoke an AWS Step Functions state machine when a new message is added to the SQS queue
- N. Use the Step Functions state machine to delete the image tag for images that have Critical or High severity finding
- O. Notify the development team by using Amazon Simple Email Service (Amazon SES).

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonECR/latest/userguide/ecr-eventbridge.html> "Activating an AWS Step Functions state machine"

<https://docs.aws.amazon.com/step-functions/latest/dg/tutorial-creating-lambda-state-machine.html>

NEW QUESTION 77

- (Exam Topic 2)

A company operates a proxy server on a fleet of Amazon EC2 instances. Partners in different countries use the proxy server to test the company's functionality. The EC2 instances are running in a VPC, and the instances have access to the internet.

The company's security policy requires that partners can access resources only from domains that the company owns.

Which solution will meet these requirements?

- A. Create an Amazon Route 53 Resolver DNS Firewall domain list that contains the allowed domains. Configure a DNS Firewall rule group with a rule that has a high numeric value that blocks all requests
- B. Configure a rule that has a low numeric value that allows requests for domains in the allowed list
- C. Associate the rule group with the VPC.
- D. Create an Amazon Route 53 Resolver DNS Firewall domain list that contains the allowed domains. Configure a Route 53 outbound endpoint
- E. Associate the outbound endpoint with the VPC
- F. Associate the domain list with the outbound endpoint.
- G. Create an Amazon Route 53 traffic flow policy to match the allowed domain
- H. Configure the traffic flow policy to forward requests that match to the Route 53 Resolver
- I. Associate the traffic flow policy with the VPC.
- J. Create an Amazon Route 53 outbound endpoint
- K. Associate the outbound endpoint with the VPC
- L. Configure a Route 53 traffic flow policy to forward requests for allowed domains to the outbound endpoint
- M. Associate the traffic flow policy with the VPC.

Answer: A

Explanation:

The company should create an Amazon Route 53 Resolver DNS Firewall domain list that contains the allowed domains. The company should configure a DNS Firewall rule group with a rule that has a high numeric value that blocks all requests. The company should configure a rule that has a low numeric value that allows requests for domains in the allowed list. The company should associate the rule group with the VPC. This solution will meet the requirements because Amazon Route 53 Resolver DNS Firewall is a feature that enables you to filter and regulate outbound DNS traffic for your VPC. You can create reusable collections of filtering rules in DNS Firewall rule groups and associate them with your VPCs. You can specify lists of domain names to allow or block, and you can customize the responses for the DNS queries that you block. By creating a domain list with the allowed domains and a rule group with rules to allow or block requests based on the domain list, the company can enforce its security policy and control access to sites.

The other options are not correct because:

- Configuring a Route 53 outbound endpoint and associating it with the VPC would not help with filtering outbound DNS traffic. A Route 53 outbound endpoint is a resource that enables you to forward DNS queries from your VPC to your network over AWS Direct Connect or VPN connections. It does not provide any filtering capabilities.
- Creating a Route 53 traffic flow policy to match the allowed domains would not help with filtering outbound DNS traffic. A Route 53 traffic flow policy is a resource that enables you to route traffic based on multiple criteria, such as endpoint health, geographic location, and latency. It does not provide any filtering capabilities.

➤ Creating a Gateway Load Balancer (GWLB) would not help with filtering outbound DNS traffic. A GWLB is a service that enables you to deploy, scale, and manage third-party virtual appliances such as firewalls, intrusion detection and prevention systems, and deep packet inspection systems in the cloud⁴. It does not provide any filtering capabilities.

References:

- <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver-dns-firewall.html>
- <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver-outbound-endpoints.html>
- <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/traffic-flow.html>
- <https://docs.aws.amazon.com/elasticloadbalancing/latest/gateway/introduction.html>

NEW QUESTION 78

- (Exam Topic 2)

A company is deploying a new web-based application and needs a storage solution for the Linux application servers. The company wants to create a single location for updates to application data for all instances. The active dataset will be up to 100 GB in size. A solutions architect has determined that peak operations will occur for 3 hours daily and will require a total of 225 MiBps of read throughput.

The solutions architect must design a Multi-AZ solution that makes a copy of the data available in another AWS Region for disaster recovery (DR). The DR copy has an RPO of less than 1 hour.

Which solution will meet these requirements?

- A. Deploy a new Amazon Elastic File System (Amazon EFS) Multi-AZ file system
- B. Configure the file system for 75 MiBps of provisioned throughput
- C. Implement replication to a file system in the DR Region.
- D. Deploy a new Amazon FSx for Lustre file system
- E. Configure Bursting Throughput mode for the file system
- F. Use AWS Backup to back up the file system to the DR Region.
- G. Deploy a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume with 225 MiBps of throughput
- H. Enable Multi-Attach for the EBS volume
- I. Use AWS Elastic Disaster Recovery to replicate the EBS volume to the DR Region.
- J. Deploy an Amazon FSx for OpenZFS file system in both the production Region and the DR Region. Create an AWS DataSync scheduled task to replicate the data from the production file system to the DR file system every 10 minutes.

Answer: A

Explanation:

The company should deploy a new Amazon Elastic File System (Amazon EFS) Multi-AZ file system. The company should configure the file system for 75 MiBps of provisioned throughput. The company should implement replication to a file system in the DR Region. This solution will meet the requirements because Amazon EFS is a serverless, fully elastic file storage service that lets you share file data without provisioning or managing storage capacity and performance. Amazon EFS is built to scale on demand to petabytes without disrupting applications, growing and shrinking automatically as you add and remove files¹. By deploying a new Amazon EFS Multi-AZ file system, the company can create a single location for updates to application data for all instances. A Multi-AZ file system replicates data across multiple Availability Zones (AZs) within a Region, providing high availability and durability². By configuring the file system for 75 MiBps of provisioned throughput, the company can ensure that it meets the peak operations requirement of 225 MiBps of read throughput. Provisioned throughput is a feature that enables you to specify a level of throughput that the file system can drive independent of the file system's size or burst credit balance³. By implementing replication to a file system in the DR Region, the company can make a copy of the data available in another AWS Region for disaster recovery. Replication is a feature that enables you to replicate data from one EFS file system to another EFS file system across AWS Regions. The replication process has an RPO of less than 1 hour.

The other options are not correct because:

- Deploying a new Amazon FSx for Lustre file system would not provide a single location for updates to application data for all instances. Amazon FSx for Lustre is a fully managed service that provides cost-effective, high-performance storage for compute workloads. However, it does not support concurrent write access from multiple instances. Using AWS Backup to back up the file system to the DR Region would not provide real-time replication of data. AWS Backup is a service that enables you to centralize and automate data protection across AWS services. However, it does not support continuous data replication or cross-Region disaster recovery.
- Deploying a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume with 225 MiBps of throughput would not provide a single location for updates to application data for all instances. Amazon EBS is a service that provides persistent block storage volumes for use with Amazon EC2 instances. However, it does not support concurrent access from multiple instances, unless Multi-Attach is enabled. Enabling Multi-Attach for the EBS volume would not provide Multi-AZ resilience or cross-Region replication. Multi-Attach is a feature that enables you to attach an EBS volume to multiple EC2 instances within the same Availability Zone. Using AWS Elastic Disaster Recovery to replicate the EBS volume to the DR Region would not provide real-time replication of data. AWS Elastic Disaster Recovery (AWS DRS) is a service that enables you to orchestrate and automate disaster recovery workflows across AWS Regions. However, it does not support continuous data replication or sub-hour RPOs.
- Deploying an Amazon FSx for OpenZFS file system in both the production Region and the DR Region would not be as simple or cost-effective as using Amazon EFS. Amazon FSx for OpenZFS is a fully managed service that provides high-performance storage with strong data consistency and advanced data management features for Linux workloads. However, it requires more configuration and management than Amazon EFS, which is serverless and fully elastic. Creating an AWS DataSync scheduled task to replicate the data from the production file system to the DR file system every 10 minutes would not provide real-time replication of data. AWS DataSync is a service that enables you to transfer data between on-premises storage and AWS services, or between AWS services. However, it does not support continuous data replication or sub-minute RPOs.

References:

- <https://aws.amazon.com/efs/>
- <https://docs.aws.amazon.com/efs/latest/ug/how-it-works.html#how-it-works-azs>
- <https://docs.aws.amazon.com/efs/latest/ug/performance.html#provisioned-throughput>
- <https://docs.aws.amazon.com/efs/latest/ug/replication.html>
- <https://aws.amazon.com/fsx/lustre/>
- <https://aws.amazon.com/backup/>
- <https://aws.amazon.com/ebs/>
- <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volumes-multi.html>

NEW QUESTION 82

- (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 Director
- 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 84

- (Exam Topic 2)

A company has millions of objects in an Amazon S3 bucket. The objects are in the S3 Standard storage class. All the S3 objects are accessed frequently. The number of users and applications that access the objects is increasing rapidly. The objects are encrypted with server-side encryption with AWS KMS Keys (SSE-KMS).

A solutions architect reviews the company's monthly AWS invoice and notices that AWS KMS costs are increasing because of the high number of requests from Amazon S3. The solutions architect needs to optimize costs with minimal changes to the application.

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

- A. Create a new S3 bucket that has server-side encryption with customer-provided keys (SSE-C) as the encryption type
- B. Copy the existing objects to the new S3 bucket
- C. Specify SSE-C.
- D. Create a new S3 bucket that has server-side encryption with Amazon S3 managed keys (SSE-S3) as the encryption type
- E. Use S3 Batch Operations to copy the existing objects to the new S3 bucket
- F. Specify SSE-S3.
- G. Use AWS CloudHSM to store the encryption key
- H. Create a new S3 bucket
- I. Use S3 Batch Operations to copy the existing objects to the new S3 bucket
- J. Encrypt the objects by using the keys from CloudHSM.
- K. Use the S3 Intelligent-Tiering storage class for the S3 bucket
- L. Create an S3 Intelligent-Tiering archive configuration to transition objects that are not accessed for 90 days to S3 Glacier Deep Archive.

Answer: B

Explanation:

To reduce the volume of Amazon S3 calls to AWS KMS, use Amazon S3 bucket keys, which are protected encryption keys that are reused for a limited time in Amazon S3. Bucket keys can reduce costs for AWS KMS requests by up to 99%. You can configure a bucket key for all objects in an Amazon S3 bucket, or for a specific object in an Amazon S3 bucket. https://docs.aws.amazon.com/fr_fr/kms/latest/developerguide/services-s3.html

NEW QUESTION 86

- (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 88

- (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 90

- (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 91

- (Exam Topic 2)

A solutions architect wants to cost-optimize and appropriately size Amazon EC2 instances in a single AWS account. The solutions architect wants to ensure that the instances are optimized based on CPU, memory, and network metrics.

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

- A. Purchase AWS Business Support or AWS Enterprise Support for the account.
- B. Turn on AWS Trusted Advisor and review any "Low Utilization Amazon EC2 Instances" recommendations.
- C. Install the Amazon CloudWatch agent and configure memory metric collection on the EC2 instances.
- D. Configure AWS Compute Optimizer in the AWS account to receive findings and optimization recommendations.
- E. Create an EC2 Instance Savings Plan for the AWS Regions, instance families, and operating systems of interest.

Answer: BD

Explanation:

AWS Trusted Advisor is a service that provides real-time guidance to help users provision their resources following AWS best practices¹. One of the Trusted Advisor checks is "Low Utilization Amazon EC2 Instances", which identifies EC2 instances that appear to be underutilized based on CPU, network I/O, and disk I/O metrics¹. This check can help users optimize the cost and size of their EC2 instances by recommending smaller or more appropriate instance types.

AWS Compute Optimizer is a service that analyzes the configuration and utilization metrics of AWS resources and generates optimization recommendations to reduce the cost and improve the performance of workloads². Compute Optimizer supports four types of AWS resources: EC2 instances, EBS volumes, ECS services on AWS Fargate, and Lambda functions². For EC2 instances, Compute Optimizer evaluates the vCPUs, memory, storage, and other specifications, as well as the CPU utilization, network in and out, disk read and write, and other utilization metrics of currently running instances³. It then recommends optimal instance types based on price-performance trade-offs.

Option A is incorrect because purchasing AWS Business Support or AWS Enterprise Support for the account will not directly help with cost-optimization and sizing

of EC2 instances. However, these support plans do provide access to more Trusted Advisor checks than the basic support plan¹. Option C is incorrect because installing the Amazon CloudWatch agent and configuring memory metric collection on the EC2 instances will not provide any optimization recommendations by itself. However, memory metrics can be used by Compute Optimizer to enhance its recommendations if enabled³. Option E is incorrect because creating an EC2 Instance Savings Plan for the AWS Regions, instance families, and operating systems of interest will not help with cost-optimization and sizing of EC2 instances. Savings Plans are a flexible pricing model that offer lower prices on Amazon EC2 usage in exchange for a commitment to a consistent amount of usage for a 1- or 3-year term⁴. Savings Plans do not affect the configuration or utilization of EC2 instances.

NEW QUESTION 92

- (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 93

- (Exam Topic 2)

A solutions architect must provide a secure way for a team of cloud engineers to use the AWS CLI to upload objects into an Amazon S3 bucket. Each cloud engineer has an IAM user, IAM access keys, and a virtual multi-factor authentication (MFA) device. The IAM users for the cloud engineers are in a group that is named S3-access. The cloud engineers must use MFA to perform any actions in Amazon S3. Which solution will meet these requirements?

- A. Attach a policy to the S3 bucket to prompt the IAM user for an MFA code when the IAM user performs actions on the S3 bucket. Use IAM access keys with the AWS CLI to call Amazon S3.
- B. Update the trust policy for the S3-access group to require principals to use MFA when principals assume the group. Use IAM access keys with the AWS CLI to call Amazon S3.
- C. Attach a policy to the S3-access group to deny all S3 actions unless MFA is present. Use IAM access keys with the AWS CLI to call Amazon S3.
- D. Attach a policy to the S3-access group to deny all S3 actions unless MFA is present. Request temporary credentials from AWS Security Token Service (AWS STS). Attach the temporary credentials in a profile that Amazon S3 will reference when the user performs actions in Amazon S3.

Answer: D

Explanation:

The company should attach a policy to the S3-access group to deny all S3 actions unless MFA is present. The company should request temporary credentials from AWS Security Token Service (AWS STS). The company should attach the temporary credentials in a profile that Amazon S3 will reference when the user performs actions in Amazon S3. This solution will meet the requirements because AWS STS is a service that enables you to request temporary, limited-privilege

credentials for IAM users or for users that you authenticate (federated users). You can use MFA with AWS STS to provide an extra layer of security when requesting temporary credentials¹. You can use the `sts get-session-token` AWS CLI command to request temporary credentials that include an MFA token². You can then use these credentials with the AWS CLI to access Amazon S3 resources. To do this, you need to attach a policy to the IAM group that denies all S3 actions unless MFA is present³. You also need to create a profile in the AWS CLI configuration file that references the temporary credentials.

The other options are not correct because:

- Attaching a policy to the S3 bucket to prompt the IAM user for an MFA code when the IAM user performs actions on the S3 bucket would not work because policies attached to S3 buckets cannot enforce MFA authentication. Policies attached to S3 buckets are resource-based policies that define what actions can be performed on the bucket and by whom. They do not have any logic to prompt for an MFA code or verify it.
- Updating the trust policy for the S3-access group to require principals to use MFA when principals assume the group would not work because trust policies are used for roles, not groups. Trust policies are policies that define which principals can assume a role. They do not apply to groups, which are collections of IAM users that share permissions.
- Creating an Amazon Route 53 Resolver DNS Firewall domain list that contains the allowed domains and configuring a DNS Firewall rule group with rules to allow or block requests based on the domain list would not help with enforcing MFA authentication for Amazon S3 actions. Amazon Route 53 Resolver DNS Firewall is a feature that enables you to filter and regulate outbound DNS traffic for your VPC. You can create reusable collections of filtering rules in DNS Firewall rule groups and associate them with your VPCs. You can specify lists of domain names to allow or block, and you can customize the responses for the DNS queries that you block. This feature is useful for controlling access to sites and blocking DNS-level threats, but not for requiring MFA authentication.

References:

- https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_temp.html
- https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_mfa_enable_cliapi.html
- https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_mfa_sample-policies.html
- <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-profiles.html>
- <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver-dns-firewall.html>

NEW QUESTION 95

- (Exam Topic 2)

A company runs its sales reporting application in an AWS Region in the United States. The application uses an Amazon API Gateway Regional API and AWS Lambda functions to generate on-demand reports from data in an Amazon RDS for MySQL database. The frontend of the application is hosted on Amazon S3 and is accessed by users through an Amazon CloudFront distribution. The company is using Amazon Route 53 as the DNS service for the domain. Route 53 is configured with a simple routing policy to route traffic to the API Gateway API.

In the next 6 months, the company plans to expand operations to Europe. More than 90% of the database traffic is read-only traffic. The company has already deployed an API Gateway API and Lambda functions in the new Region.

A solutions architect must design a solution that minimizes latency for users who download reports. Which solution will meet these requirements?

- A. Use an AWS Database Migration Service (AWS DMS) task with full load to replicate the primary database in the original Region to the database in the new Region
- B. Change the Route 53 record to latency-based routing to connect to the API Gateway API.
- C. Use an AWS Database Migration Service (AWS DMS) task with full load plus change data capture (CDC) to replicate the primary database in the original Region to the database in the new Region
- D. Change the Route 53 record to geolocation routing to connect to the API Gateway API.
- E. Configure a cross-Region read replica for the RDS database in the new Region
- F. Change the Route 53 record to latency-based routing to connect to the API Gateway API.
- G. Configure a cross-Region read replica for the RDS database in the new Region
- H. Change the Route 53 record to geolocation routing to connect to the API

Answer: C

Explanation:

The company should configure a cross-Region read replica for the RDS database in the new Region. The company should change the Route 53 record to latency-based routing to connect to the API Gateway API. This solution will meet the requirements because a cross-Region read replica is a feature that enables you to create a MariaDB, MySQL, Oracle, PostgreSQL, or SQL Server read replica in a different Region from the source DB instance. You can use cross-Region read replicas to improve availability and disaster recovery, scale out globally, or migrate an existing database to a new Region¹. By creating a cross-Region read replica for the RDS database in the new Region, the company can have a standby copy of its primary database that can serve read-only traffic from users in Europe. A latency-based routing policy is a feature that enables you to route traffic based on the latency between your users and your resources. You can use latency-based routing to route traffic to the resource that provides the best latency². By changing the Route 53 record to latency-based routing, the company can minimize latency for users who download reports by connecting them to the API Gateway API in the Region that provides the best response time.

The other options are not correct because:

- Using AWS Database Migration Service (AWS DMS) to replicate the primary database in the original Region to the database in the new Region would not be as cost-effective or simple as using a cross-Region read replica. AWS DMS is a service that enables you to migrate relational databases, data warehouses, NoSQL databases, and other types of data stores. You can use AWS DMS to perform one-time migrations or continuous data replication with high availability and consolidate databases into a petabyte-scale data warehouse³. However, AWS DMS requires more configuration and management than creating a cross-Region read replica, which is fully managed by Amazon RDS. AWS DMS also incurs additional charges for replication instances and tasks.
- Creating an Amazon API Gateway Data API service integration with Amazon Redshift would not help with disaster recovery or minimizing latency. 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 from an RDS database.
- Creating an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster would not help with disaster recovery or minimizing latency. 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 from an RDS database.

References:

- <https://aws.amazon.com/dms/>
- <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 99

- (Exam Topic 2)

A company runs a processing engine in the AWS Cloud. The engine processes environmental data from logistics centers to calculate a sustainability index. The company has millions of devices in logistics centers that are spread across Europe. The devices send information to the processing engine through a RESTful API. The API experiences unpredictable bursts of traffic. The company must implement a solution to process all data that the devices send to the processing engine. Data loss is unacceptable.

Which solution will meet these requirements?

- A. Create an Application Load Balancer (ALB) for the RESTful API. Create an Amazon Simple Queue Service (Amazon SQS) queue. Create a listener and a target group for the ALB. Add the SQS queue as the target. Use a container that runs in Amazon Elastic Container Service (Amazon ECS) with the Fargate launch type to process messages in the queue.
- B. Create an Amazon API Gateway HTTP API that implements the RESTful API. Create an Amazon Simple Queue Service (Amazon SQS) queue. Create an API Gateway service integration with the SQS queue. Create an AWS Lambda function to process messages in the SQS queue.
- C. Create an Amazon API Gateway REST API that implements the RESTful API. Create a fleet of Amazon EC2 instances in an Auto Scaling group. Create an API Gateway Auto Scaling group proxy integration. Use the EC2 instances to process incoming data.
- D. Create an Amazon CloudFront distribution for the RESTful API. Create a data stream in Amazon Kinesis Data Streams. Set the data stream as the origin for the distribution. Create an AWS Lambda function to consume and process data in the data stream.

Answer: A

Explanation:

It will use the ALB to handle the unpredictable bursts of traffic and route it to the SQS queue. The SQS queue will act as a buffer to store incoming data temporarily, and the container running in Amazon ECS with the Fargate launch type will process messages in the queue. This approach will ensure that all data is processed and prevent data loss.

NEW QUESTION 102

.....

Relate Links

100% Pass Your AWS-Certified-Solutions-Architect-Professional Exam with ExamBible Prep Materials

<https://www.exambible.com/AWS-Certified-Solutions-Architect-Professional-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>