



**Amazon**

## **Exam Questions AWS-Certified-Solutions-Architect-Professional**

Amazon AWS Certified Solutions Architect Professional

**NEW QUESTION 1**

- (Exam Topic 2)

A retail company needs to provide a series of data files to another company, which is its business partner. These files are saved in an Amazon S3 bucket under Account A, which belongs to the retail company. The business partner company wants one of its IAM users, User\_DataProcessor, to access the files from its own AWS account (Account B).

Which combination of steps must the companies take so that User\_DataProcessor can access the S3 bucket successfully? (Select TWO.)

- A. Turn on the cross-origin resource sharing (CORS) feature for the S3 bucket in Account A.
- B. In Account A, set the S3 bucket policy to the following:

```
{
  "Effect": "Allow",
  "Action": [
    "s3:GetObject",
    "s3:ListBucket"
  ],
  "Resource": "arn:aws:s3:::AccountABucketName/*"
}
```

- D. In Account A, set the S3 bucket policy to the following:

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "arn:aws:iam::AccountB:user/User_DataProcessor"
  },
  "Action": [
    "s3:GetObject",
    "s3:ListBucket"
  ],
  "Resource": [
    "arn:aws:s3:::AccountABucketName/*"
  ]
}
```

- F. In Account B, set the permissions of User\_DataProcessor to the following:

```
{
  "Effect": "Allow",
  "Action": [
    "s3:GetObject",
    "s3:ListBucket"
  ],
  "Resource": "arn:aws:s3:::AccountABucketName/*"
}
```

- H. In Account B, set the permissions of User\_DataProcessor to the following:

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "arn:aws:iam::AccountB:user/User_DataProcessor"
  },
  "Action": [
    "s3:GetObject",
    "s3:ListBucket"
  ],
  "Resource": [
    "arn:aws:s3:::AccountABucketName/*"
  ]
}
```

**Answer:** CD

**Explanation:**

<https://aws.amazon.com/premiumsupport/knowledge-center/cross-account-access-s3/>

**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-month<sup>1</sup>. 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 days<sup>2</sup>, otherwise they incur a pro-rated charge equal to the storage charge for the remaining days<sup>1</sup>. 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 has VPC flow logs enabled for its NAT gateway. The company is seeing Action = ACCEPT for inbound traffic that comes from public IP address 198.51.100.2 destined for a private Amazon EC2 instance.

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

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 5**

- (Exam Topic 2)

A company uses a load balancer to distribute traffic to Amazon EC2 instances in a single Availability Zone. The company is concerned about security and wants a solutions architect to re-architect the solution to meet the following requirements:

- Inbound requests must be filtered for common vulnerability attacks.
- Rejected requests must be sent to a third-party auditing application.
- All resources should be highly available. Which solution meets these requirements?

- A. Configure a Multi-AZ Auto Scaling group using the application's AM
- B. Create an Application Load Balancer (ALB) and select the previously created Auto Scaling group as the target
- C. Use Amazon Inspector to monitor traffic to the ALB and EC2 instance
- D. Create a web ACL in WA
- E. Create an AWS WAF using the web ACL and AL

- F. Use an AWS Lambda function to frequently push the Amazon Inspector report to the third-party auditing application.
- G. Configure an Application Load Balancer (ALB) and add the EC2 instances as targets Create a web ACL in WA
- H. Create an AWS WAF using the web ACL and ALB name and enable logging with Amazon CloudWatch Log
- I. Use an AWS Lambda function to frequently push the logs to the third-party auditing application.
- J. Configure an Application Load Balancer (ALB) along with a target group adding the EC2 instances as target
- K. Create an Amazon Kinesis Data Firehose with the destination of the third-party auditing applicatio
- L. Create a web ACL in WA
- M. Create an AWS WAF using the web ACL and ALB then enable logging by selecting the Kinesis Data Firehose as the destinatio
- N. Subscribe to AWS Managed Rules in AWS Marketplace, choosing the WAF as the subscriber.
- O. Configure a Multi-AZ Auto Scaling group using the application's AM
- P. Create an Application Load Balancer (ALB) and select the previously created Auto Scaling group as the targe
- Q. Create an Amazon Kinesis Data Firehose with a destination of the third-party auditing applicatio
- R. Create a web ACL inWA
- S. Create an AWS WAF using the WebACL and ALB then enable logging by selecting the Kinesis Data Firehose as the destinatio
- T. Subscribe to AWS Managed Rules in AWS Marketplace, choosing the WAF as the subscriber.

**Answer:** D

**Explanation:**

<https://docs.aws.amazon.com/waf/latest/developerguide/marketplace-managed-rule-groups.html>

**NEW QUESTION 6**

- (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 7**

- (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 theCloudFormation 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 8**

- (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
- U. Use an Amazon Route 53 failover routing policy to automatically fail over to the DR Region in the event of a disaster.

**Answer: B**

**Explanation:**

The company should use infrastructure as code (IaC) to provision the new infrastructure in the DR Region. The company should create a cross-Region read replica for the DB instance. The company should set up AWS Elastic Disaster Recovery to continuously replicate the EC2 instances to the DR Region. The company should run the EC2 instances at the minimum capacity in the DR Region. The company should use an Amazon Route 53 failover routing policy to automatically fail over to the DR Region in the event of a disaster. The company should increase the desired capacity of the Auto Scaling group. This solution will meet the requirements most cost-effectively because AWS Elastic Disaster Recovery (AWS DRS) is a service that minimizes downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications using affordable storage, minimal compute, and point-in-time recovery. AWS DRS enables RPOs of seconds and RTOs of minutes. AWS DRS continuously replicates data from the source servers to a staging area subnet in the DR Region, where it uses low-cost storage and minimal compute resources to maintain ongoing replication. In the event of a disaster, AWS DRS automatically converts the servers to boot and run natively on AWS and launches recovery instances on AWS within minutes. By using AWS DRS, the company can save costs by removing idle recovery site resources and paying for the full disaster recovery site only when needed. By creating a cross-Region read replica for the DB instance, the company can have a standby copy of its primary database in a different AWS Region. By using infrastructure as code (IaC), the company can provision the new infrastructure in the DR Region in an automated and consistent way. 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 9**

- (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). 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 bucket. Configuring redirection of HTTP requests to HTTPS requests in CloudFront will enable encrypting the data in transit using SSL/TLS.

**NEW QUESTION 10**

- (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 10**

- (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](https://docs.aws.amazon.com/organizations/latest/APIReference/API_RemoveAccountFromOrganization.html)

**NEW QUESTION 15**

- (Exam Topic 2)

A company is running an application in the AWS Cloud. The application collects and stores a large amount of unstructured data in an Amazon S3 bucket. The S3 bucket contains several terabytes of data and uses the S3 Standard storage class. The data increases in size by several gigabytes every day.

The company needs to query and analyze the data. The company does not access data that is more than 1 year old. However, the company must retain all the data indefinitely for compliance reasons.

Which solution will meet these requirements MOST cost-effectively?

- A. Use S3 Select to query the data.
- B. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Glacier Deep Archive.
- C. Use Amazon Redshift Spectrum to query the data.

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

**Answer: C**

**Explanation:**

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

**NEW QUESTION 17**

- (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 21**

- (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. AWS Secrets Manager helps you to securely encrypt, store, and retrieve credentials for your databases and other services. Attaching a role to each Lambda function to provide access to the secret will enable retrieving the credentials programmatically. 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.

**NEW QUESTION 22**

- (Exam Topic 2)

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

- A. Create a stack set in the Organizations member account
- B. Use service-managed permission
- C. Set deployment options to deploy to an organization
- D. Use CloudFormation StackSets drift detection.

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 23**

- (Exam Topic 2)

A company is updating an application that customers use to make online orders. The number of attacks on the application by bad actors has increased recently. The company will host the updated application on an Amazon Elastic Container Service (Amazon ECS) cluster. The company will use Amazon DynamoDB to store application data. A public Application Load Balancer (ALB) will provide end users with access to the application. The company must prevent attacks and ensure business continuity with minimal service interruptions during an ongoing attack.

Which combination of steps will meet these requirements MOST cost-effectively? (Select TWO.)

- A. Create an Amazon CloudFront distribution with the ALB as the origin
- B. Add a custom header and random value on the CloudFront domain
- C. Configure the ALB to conditionally forward traffic if the header and value match.
- D. Deploy the application in two AWS Region
- E. Configure Amazon Route 53 to route to both Regions with equal weight.
- F. Configure auto scaling for Amazon ECS task
- G. Create a DynamoDB Accelerator (DAX) cluster.
- H. Configure Amazon ElastiCache to reduce overhead on DynamoDB.
- I. Deploy an AWS WAF web ACL that includes an appropriate rule group
- J. Associate the web ACL with the Amazon CloudFront distribution.

**Answer:** AE

**Explanation:**

The company should create an Amazon CloudFront distribution with the ALB as the origin. The company should add a custom header and random value on the CloudFront domain. The company should configure the ALB to conditionally forward traffic if the header and value match. The company should also deploy an AWS WAF web ACL that includes an appropriate rule group. The company should associate the web ACL with the Amazon CloudFront distribution. This solution will meet the requirements most cost-effectively because Amazon CloudFront is a fast content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency, high transfer speeds, all within a developer-friendly environment<sup>1</sup>. By creating an Amazon CloudFront distribution with the ALB as the origin, the company can improve the performance and availability of its application by caching static content at edge locations closer to end users. By adding a custom header and random value on the CloudFront domain, the company can prevent direct access to the ALB and ensure that only requests from CloudFront are forwarded to the ECS tasks. By configuring the ALB to conditionally forward traffic if the header and value match, the company can implement origin access identity (OAI) for its ALB origin. OAI is a feature that enables you to restrict access to your content by requiring users to access your content through CloudFront URLs<sup>2</sup>. By deploying an AWS WAF web ACL that includes an appropriate rule group, the company can prevent attacks and ensure business continuity with minimal service interruptions during an ongoing attack. AWS WAF is a web application firewall that lets you monitor and control web requests that are forwarded to your web applications. You can use AWS WAF to define customizable web security rules that control which traffic can access your web applications and which traffic should be blocked<sup>3</sup>. By associating the web ACL with the Amazon CloudFront distribution, the company can apply the web security rules to all requests that are forwarded by CloudFront.

The other options are not correct because:

- Deploying the application in two AWS Regions and configuring Amazon Route 53 to route to both Regions with equal weight would not prevent attacks or ensure business continuity. Amazon Route 53 is a highly available and scalable cloud Domain Name System (DNS) web service that routes end users to Internet applications by translating names like `www.example.com` into numeric IP addresses<sup>4</sup>. However, routing traffic to multiple Regions would not protect against attacks or provide failover in case of an outage. It would also increase operational complexity and costs compared to using CloudFront and AWS WAF.
- Configuring auto scaling for Amazon ECS tasks and creating a DynamoDB Accelerator (DAX) cluster would not prevent attacks or ensure business continuity. Auto scaling is a feature that enables you to automatically adjust your ECS tasks based on demand or a schedule. DynamoDB Accelerator (DAX) is a fully managed, highly available, in-memory cache for DynamoDB that delivers up to a 10x performance improvement. However, these features would not protect against attacks or provide failover in case of an outage. They would also increase operational complexity and costs compared to using CloudFront and AWS WAF.
- Configuring Amazon ElastiCache to reduce overhead on DynamoDB would not prevent attacks or ensure business continuity. Amazon ElastiCache is a fully managed in-memory data store service that makes it easy to deploy, operate, and scale popular open-source compatible in-memory data stores. However, this service would not protect against attacks or provide failover in case of an outage. It would also increase operational complexity and costs compared to using CloudFront and AWS WAF.

References:

- <https://aws.amazon.com/cloudfront/>
- <https://aws.amazon.com/waf/>
- <https://aws.amazon.com/route53/>
- <https://aws.amazon.com/dynamodb/dax/>
- <https://aws.amazon.com/elasticache/>

**NEW QUESTION 26**

- (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 capacity
- 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 capacity
- 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 28**

- (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 32**

- (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 37**

- (Exam Topic 2)

A company has migrated a legacy application to the AWS Cloud. The application runs on three Amazon EC2 instances that are spread across three Availability Zones. One EC2 instance is in each Availability Zone. The EC2 instances are running in three private subnets of the VPC and are set up as targets for an Application Load Balancer (ALB) that is associated with three public subnets.

The application needs to communicate with on-premises systems. Only traffic from IP addresses in the company's IP address range are allowed to access the on-premises systems. The company's security team is bringing only one IP address from its internal IP address range to the cloud. The company has added this IP address to the allow list for the company firewall. The company also has created an Elastic IP address for this IP address.

A solutions architect needs to create a solution that gives the application the ability to communicate with the on-premises systems. The solution also must be able to mitigate failures automatically.

Which solution will meet these requirements?

- A. Deploy three NAT gateways, one in each public subne
- B. Assign the Elastic IP address to the NAT gateway
- C. Turn on health checks for the NAT gateway
- D. If a NAT gateway fails a health check, recreate the NAT gateway and assign the Elastic IP address to the new NAT gateway.
- E. Replace the ALB with a Network Load Balancer (NLB). Assign the Elastic IP address to the NLB Turn on health checks for the NL
- F. In the case of a failed health check, redeploy the NLB in different subnets.
- G. Deploy a single NAT gateway in a public subne
- H. Assign the Elastic IP address to the NAT gateway.Use Amazon CloudWatch with a custom metric tomonitor the NAT gatewa
- I. If the NAT gateway is unhealthy, invoke an AWS Lambda function to create a new NAT gateway in a different subne
- J. Assign the Elastic IP address to the new NAT gateway.
- K. Assign the Elastic IP address to the AL
- L. Create an Amazon Route 53 simple record with the Elastic IP address as the valu
- M. Create a Route 53 health chec
- N. In the case of a failed health check, recreate the ALB in different subnets.

**Answer: C**

**Explanation:**

to connect out from the private subnet you need an NAT gateway and since only one Elastic IP whitelisted on firewall its one NATGateway at time and if AZ failure happens Lambda creates a new NATGATEWAY in a different AZ using the Same Elastic IP ,dont be tempted to select D since application that needs to connect is on a private subnet whose outbound connections use the NATGateway Elastic IP

**NEW QUESTION 38**

- (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 request
- B. Configure a rule that has a low numeric value that allows requests for domains in the allowed lis
- 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 VP
- 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 Resolve
- I. Associate the traffic flow policy with the VPC.
- J. Create an Amazon Route 53 outbound endpoint
- K. Associate the outbound endpoint with the VP
- 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 block1. 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 connections2. 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 latency3. 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 41

- (Exam Topic 2)

A company is running a web application in a VPC. The web application runs on a group of Amazon EC2 instances behind an Application Load Balancer (ALB). The ALB is using AWS WAF.

An external customer needs to connect to the web application. The company must provide IP addresses to all external customers.

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

- A. Replace the ALB with a Network Load Balancer (NLB). Assign an Elastic IP address to the NLB.
- B. Allocate an Elastic IP address
- C. Assign the Elastic IP address to the ALB. Provide the Elastic IP address to the customer.
- D. Create an AWS Global Accelerator standard accelerator
- E. Specify the ALB as the accelerator's endpoint. Provide the accelerator's IP addresses to the customer.
- F. Configure an Amazon CloudFront distribution
- G. Set the ALB as the origin
- H. Ping the distribution's DNS name to determine the distribution's public IP address
- I. Provide the IP address to the customer.

**Answer: C**

#### Explanation:

<https://docs.aws.amazon.com/global-accelerator/latest/dg/about-accelerators.alb-accelerator.html> Option A is wrong. AWS WAF does not support associating with NLB.

<https://docs.aws.amazon.com/waf/latest/developerguide/waf-chapter.html> Option B is wrong. An ALB does not support an Elastic IP address.

<https://aws.amazon.com/elasticloadbalancing/features/>

#### NEW QUESTION 46

- (Exam Topic 2)

A company runs an application on AWS. The company curates data from several different sources. The company uses proprietary algorithms to perform data transformations and aggregations. After the company performs ETL processes, the company stores the results in Amazon Redshift tables. The company sells this data to other companies. The company downloads the data as files from the Amazon Redshift tables and transmits the files to several data customers by using FTP. The number of data customers has grown significantly. Management of the data customers has become difficult.

The company will use AWS Data Exchange to create a data product that the company can use to share data with customers. The company wants to confirm the identities of the customers before the company shares data.

The customers also need access to the most recent data when the company publishes the data. Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS Data Exchange for APIs to share data with customer
- B. Configure subscription verification In the AWS account of the company that produces the data, create an Amazon API Gateway Data API service integration with Amazon Redshift
- C. Require the data customers to subscribe to the data product In the AWS account of the company that produces the data, create an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster
- D. Create an Amazon API Gateway Data API service integration with Amazon Redshift
- E. Configure subscription verification
- F. Require the data customers to subscribe to the data product.
- G. Download the data from the Amazon Redshift tables to an Amazon S3 bucket periodically
- H. Use AWS Data Exchange for S3 to share data with customers.
- I. Configure subscription verification
- J. Require the data customers to subscribe to the data product Publish the Amazon Redshift data to an Open Data on AWS Data Exchange
- K. Require the customers to subscribe to the data product in AWS Data Exchange
- L. In the AWS account of the company that produces the data, attach IAM resource-based policies to the Amazon Redshift tables to allow access only to verified AWS accounts.

**Answer: C**

#### Explanation:

The company should download the data from the Amazon Redshift tables to an Amazon S3 bucket periodically and use AWS Data Exchange for S3 to share data with customers. The company should configure subscription verification and require the data customers to subscribe to the data product. This solution will meet the requirements with the least operational overhead because AWS Data Exchange for S3 is a feature that enables data subscribers to access third-party data files directly from data providers' Amazon S3 buckets. Subscribers can easily use these files for their data analysis with AWS services without needing to create or manage data copies. Data providers can easily set up AWS Data Exchange for S3 on top of their existing S3 buckets to share direct access to an entire S3 bucket or specific prefixes and S3 objects. AWS Data Exchange automatically manages subscriptions, entitlements, billing, and payment<sup>1</sup>.

The other options are not correct because:

- Using AWS Data Exchange for APIs to share data with customers would not work because AWS Data Exchange for APIs is a feature that enables data subscribers to access third-party APIs directly from data providers' AWS accounts. Subscribers can easily use these APIs for their data analysis with AWS services without needing to manage API keys or tokens. Data providers can easily set up AWS Data Exchange for APIs on top of their existing API Gateway resources to share direct access to an entire API or specific routes and stages<sup>2</sup>. However, this feature is not suitable for sharing data from Amazon Redshift tables, which are not exposed as APIs.
- Creating an Amazon API Gateway Data API service integration with Amazon Redshift would not work because 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<sup>3</sup>. It is useful for building applications that interact with Amazon Redshift, but not for sharing data files with customers.
- Creating an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster would not work because AWS Data Exchange does not support datashares for Amazon Redshift clusters. 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<sup>4</sup>. It is useful for sharing query results and views with other users, but not for

sharing data files with customers.

➤ Publishing the Amazon Redshift data to an Open Data on AWS Data Exchange would not work because Open Data on AWS Data Exchange is a feature that enables you to find and use free and public datasets from AWS customers and partners. It is useful for accessing open and free data, but not for confirming the identities of the customers or charging them for the data.

References:

- <https://aws.amazon.com/data-exchange/why-aws-data-exchange/s3/>
- <https://aws.amazon.com/data-exchange/why-aws-data-exchange/api/>
- <https://docs.aws.amazon.com/redshift/latest/mgmt/data-api.html>
- <https://docs.aws.amazon.com/redshift/latest/dg/datashare-overview.html>
- <https://aws.amazon.com/data-exchange/open-data/>

#### NEW QUESTION 50

- (Exam Topic 2)

A company is migrating its development and production workloads to a new organization in AWS Organizations. The company has created a separate member account for development and a separate member account for production. Consolidated billing is linked to the management account. In the management account, a solutions architect needs to create an IAM user that can stop or terminate resources in both member accounts.

Which solution will meet this requirement?

- A. Create an IAM user and a cross-account role in the management account
- B. Configure the cross-account role with least privilege access to the member accounts.
- C. Create an IAM user in each member account
- D. In the management account, create a cross-account role that has least privilege access
- E. Grant the IAM users access to the cross-account role by using a trust policy.
- F. Create an IAM user in the management account
- G. In the member accounts, create an IAM group that has least privilege access
- H. Add the IAM user from the management account to each IAM group in the member accounts.
- I. Create an IAM user in the management account
- J. In the member accounts, create cross-account roles that have least privilege access
- K. Grant the IAM user access to the roles by using a trust policy.

**Answer:** D

#### Explanation:

Cross account role should be created in destination(member) account. The role has trust entity to master account.

#### NEW QUESTION 51

- (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 files<sup>1</sup>. 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 durability<sup>2</sup>. 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 balance<sup>3</sup>. 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 52

- (Exam Topic 2)

A company plans to migrate a three-tiered web application from an on-premises data center to AWS. The company developed the UI by using server-side JavaScript libraries. The business logic and API tier uses a Python-based web framework. The data tier runs on a MySQL database. The company custom built the application to meet business requirements. The company does not want to re-architect the application. The company needs a solution to replatform the application to AWS with the least possible amount of development. The solution needs to be highly available and must reduce operational overhead. Which solution will meet these requirements?

- A. Deploy the UI to a static website on Amazon S3. Use Amazon CloudFront to deliver the website. Build the business logic in a Docker image. Store the image in Amazon Elastic Container Registry (Amazon ECR). Use Amazon Elastic Container Service (Amazon ECS) with the Fargate launch type to host the website with an Application Load Balancer in front. Deploy the data layer to an Amazon Aurora MySQL DB cluster.
- B. Build the UI and business logic in Docker images. Store the images in Amazon Elastic Container Registry (Amazon ECR). Use Amazon Elastic Container Service (Amazon ECS) with the Fargate launch type to host the UI and business logic applications with an Application Load Balancer in front. Migrate the database to an Amazon RDS for MySQL Multi-AZ DB instance.
- C. Deploy the UI to a static website on Amazon S3. Use Amazon CloudFront to deliver the website. Convert the business logic to AWS Lambda functions. Integrate the functions with Amazon API Gateway. Deploy the data layer to an Amazon Aurora MySQL DB cluster.
- D. Build the UI and business logic in Docker images. Store the images in Amazon Elastic Container Registry (Amazon ECR). Use Amazon Elastic Kubernetes Service (Amazon EKS) with Fargate profiles to host the UI and business logic. Use AWS Database Migration Service (AWS DMS) to migrate the data layer to Amazon DynamoDB.

**Answer:** A

### Explanation:

This solution utilizes Amazon S3 and CloudFront to deploy the UI as a static website, which can be done with minimal development effort. The business logic and API tier can be containerized in a Docker image and stored in Amazon Elastic Container Registry (ECR) and run on Amazon Elastic Container Service (ECS) with the Fargate launch type, which allows the application to be highly available with minimal operational overhead. The data layer can be deployed on an Amazon Aurora MySQL DB cluster, which is a fully managed relational database service. Amazon Aurora provides high availability and performance for the data layer without the need for managing the underlying infrastructure.

## NEW QUESTION 56

- (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](https://docs.aws.amazon.com/fr_fr/kms/latest/developerguide/services-s3.html)

#### NEW QUESTION 57

- (Exam Topic 2)

A company runs an application on a fleet of Amazon EC2 instances that are in private subnets behind an internet-facing Application Load Balancer (ALB). The ALB is the origin for an Amazon CloudFront distribution. An AWS WAF web ACL that contains various AWS managed rules is associated with the CloudFront distribution.

The company needs a solution that will prevent internet traffic from directly accessing the ALB. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create a new web ACL that contains the same rules that the existing web ACL contain
- B. Associate the new web ACL with the ALB.
- C. Associate the existing web ACL with the ALB.
- D. Add a security group rule to the ALB to allow traffic from the AWS managed prefix list for CloudFront only.
- E. Add a security group rule to the ALB to allow only the various CloudFront IP address ranges.

**Answer:** C

#### Explanation:

<https://aws.amazon.com/about-aws/whats-new/2022/02/amazon-cloudfront-managed-prefix-list/>

#### NEW QUESTION 61

- (Exam Topic 2)

A company operates an on-premises software-as-a-service (SaaS) solution that ingests several files daily. The company provides multiple public SFTP endpoints to its customers to facilitate the file transfers. The customers add the SFTP endpoint IP addresses to their firewall allow list for outbound traffic. Changes to the SFTP endpoint IP addresses are not permitted.

The company wants to migrate the SaaS solution to AWS and decrease the operational overhead of the file transfer service.

Which solution meets these requirements?

- A. Register the customer-owned block of IP addresses in the company's AWS account
- B. Create Elastic IP addresses from the address pool and assign them to an AWS Transfer for SFTP endpoint
- C. Use AWS Transfer to store the files in Amazon S3.
- D. Add a subnet containing the customer-owned block of IP addresses to a VPC. Create Elastic IP addresses from the address pool and assign them to an Application Load Balancer (ALB). Launch EC2 instances hosting FTP services in an Auto Scaling group behind the ALB.
- E. Store the files in attached Amazon Elastic Block Store (Amazon EBS) volumes.
- F. Register the customer-owned block of IP addresses with Amazon Route 53. Create alias records in Route 53 that point to a Network Load Balancer (NLB). Launch EC2 instances hosting FTP services in an Auto Scaling group behind the NLB.
- G. Store the files in Amazon S3.
- H. Register the customer-owned block of IP addresses in the company's AWS account
- I. Create Elastic IP addresses from the address pool and assign them to an Amazon S3 VPC endpoint
- J. Enable SFTP support on the S3 bucket.

**Answer:** A

#### Explanation:

Bring your own IP addresses (BYOIP) You can bring part or all of your publicly routable IPv4 or IPv6 address range from your on-premises network to your AWS account. You continue to own the address range, but AWS advertises it on the internet by default. After you bring the address range to AWS, it appears in your AWS account as an address pool. <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-byoip.html> AWS Transfer for SFTP enables you to easily move your file transfer workloads that use the Secure Shell File Transfer Protocol (SFTP) to AWS without needing to modify your applications or manage any SFTP servers. <https://aws.amazon.com/about-aws/whats-new/2018/11/aws-transfer-for-sftp-fully-managed-sftp-for-s3/>

#### NEW QUESTION 66

- (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 69

- (Exam Topic 2)

A telecommunications company is running an application on AWS. The company has set up an AWS Direct Connect connection between the company's on-

premises data center and AWS. The company deployed the application on Amazon EC2 instances in multiple Availability Zones behind an internal Application Load Balancer (ALB). The company's clients connect from the on-premises network by using HTTPS. The TLS terminates in the ALB. The company has multiple target groups and uses path-based routing to forward requests based on the URL path. The company is planning to deploy an on-premises firewall appliance with an allow list that is based on IP address. A solutions architect must develop a solution to allow traffic flow to AWS from the on-premises network so that the clients can continue to access the application. Which solution will meet these requirements?

- A. Configure the existing ALB to use static IP addresses
- B. Assign IP addresses in multiple Availability Zones to the ALB
- C. Add the ALB IP addresses to the firewall appliance.
- D. Create a Network Load Balancer (NLB). Associate the NLB with one static IP addresses in multiple Availability Zone
- E. Create an ALB-type target group for the NLB and add the existing ALB IP addresses to the firewall appliance
- F. Update the clients to connect to the NLB.
- G. Create a Network Load Balancer (NLB). Associate the NLB with one static IP addresses in multiple Availability Zone
- H. Add the existing target groups to the NLB
- I. Update the clients to connect to the NLB
- J. Delete the ALB Add the NLB IP addresses to the firewall appliance.
- K. Create a Gateway Load Balancer (GWLB). Assign static IP addresses to the GWLB in multiple Availability Zone
- L. Create an ALB-type target group for the GWLB and add the existing ALB IP addresses to the firewall appliance
- M. Add the GWLB IP addresses to the firewall appliance
- N. Update the clients to connect to the GWLB.

**Answer: B**

**Explanation:**

The company should create a Network Load Balancer (NLB) and associate it with one static IP address in multiple Availability Zones. The company should also create an ALB-type target group for the NLB and add the existing ALB. The company should add the NLB IP addresses to the firewall appliance and update the clients to connect to the NLB. This solution will allow traffic flow to AWS from the on-premises network by using static IP addresses that can be added to the firewall appliance's allow list. The NLB will forward requests to the ALB, which will use path-based routing to forward requests to the target groups.

**NEW QUESTION 72**

- (Exam Topic 2)

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

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

**Answer: C**

**Explanation:**

The company should create a resource-based IAM policy that includes conditions for specific DynamoDB attributes (fine-grained access control). The company should attach the policy to the DynamoDB table. In the marketing team's account, the company should create an IAM role that has permissions to access the DynamoDB table in the finance team's account. This solution will meet the requirements because a resource-based IAM policy is a policy that you attach to an AWS resource (such as a DynamoDB table) to control who can access that resource and what actions they can perform on it. You can use IAM policy conditions to specify fine-grained access control for DynamoDB items and attributes. For example, you can allow or deny access to specific attributes of all items in a table by matching on attribute names<sup>1</sup>. 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 have<sup>2</sup>. 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 resources<sup>3</sup>. 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)<sup>4</sup>. 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/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/id_roles.html)
- [https://docs.aws.amazon.com/IAM/latest/UserGuide/access\\_policies\\_boundaries.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_boundaries.html)

#### NEW QUESTION 74

- (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 Region<sup>1</sup>. 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 latency<sup>2</sup>. 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 warehouse<sup>3</sup>. 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 79

- (Exam Topic 2)

A company is running a compute workload by using Amazon EC2 Spot Instances that are in an Auto Scaling group. The launch template uses two placement groups and a single instance type.

Recently, a monitoring system reported Auto Scaling instance launch failures that correlated with longer wait times for system users. The company needs to improve the overall reliability of the workload.

Which solution will meet this requirement?

- A. Replace the launch template with a launch configuration to use an Auto Scaling group that uses attribute-based instance type selection.
- B. Create a new launch template version that uses attribute-based instance type selection
- C. Configure the Auto Scaling group to use the new launch template version.
- D. Update the launch template Auto Scaling group to increase the number of placement groups.
- E. Update the launch template to use a larger instance type.

**Answer: B**

#### Explanation:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/create-asg-instance-type-requirements.html#use-attribute-based-instance-type-selection>

#### NEW QUESTION 81

- (Exam Topic 2)

A company's solutions architect is analyzing costs of a multi-application environment. The environment is deployed across multiple Availability Zones in a single AWS Region. After a recent acquisition, the company manages two organizations in AWS Organizations. The company has created multiple service provider

applications as AWS PrivateLink-powered VPC endpoint services in one organization. The company has created multiple service consumer applications in the other organization.

Data transfer charges are much higher than the company expected, and the solutions architect needs to reduce the costs. The solutions architect must recommend guidelines for developers to follow when they deploy services. These guidelines must minimize data transfer charges for the whole environment. Which guidelines meet these requirements? (Select TWO.)

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

**Answer:** CD

**Explanation:**

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

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

**NEW QUESTION 83**

- (Exam Topic 1)

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

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

What should the solutions architect recommend to meet these requirements?

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 86**

- (Exam Topic 1)

A company wants to migrate its data analytics environment from on premises to AWS. The environment consists of two simple Node.js applications. One of the applications collects sensor data and loads it into a MySQL database. The other application aggregates the data into reports. When the aggregation jobs run, some of the load jobs fail to run correctly.

The company must resolve the data loading issue. The company also needs the migration to occur without interruptions or changes for the company's customers. What should a solutions architect do to meet these requirements?

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 90**

- (Exam Topic 1)

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

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

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

Which solution will meet these requirements?

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

**Answer: C**

**Explanation:**

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

**NEW QUESTION 93**

- (Exam Topic 1)

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

What is the MOST cost-effective migration recommendation?

- A. Create a queue using Amazon SQS
- B. Configure the existing web server to publish to the new queue. When there are messages in the queue, invoke an AWS Lambda function to pull requests from the queue and process the file
- C. Store the processed files in an Amazon S3 bucket.
- D. Create a queue using Amazon SNS
- E. Configure the existing web server to publish to the new queue
- F. When there are messages in the queue, create a new Amazon EC2 instance to pull requests from the queue and process the file
- G. Store the processed files in Amazon EFS
- H. Shut down the EC2 instance after the task is complete.
- I. Create a queue using Amazon MQ
- J. Configure the existing web server to publish to the new queue. When there are messages in the queue, invoke an AWS Lambda function to pull requests from the queue and process the file
- K. Store the processed files in Amazon EFS.
- L. Create a queue using Amazon SNS
- M. Configure the existing web server to publish to the new queue
- N. Use Amazon EC2 instances in an EC2 Auto Scaling group to pull requests from the queue and process the file
- O. Scale the EC2 instances based on the SQS queue length
- P. Store the processed files in an Amazon S3 bucket.

**Answer: D**

**Explanation:**

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

**NEW QUESTION 97**

- (Exam Topic 1)

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

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

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

Simple Queue Service (Amazon SQS) queue to replace the on-premises messaging queue Configure Amazon Elastic Kubernetes Service (Amazon EKS) to host the order-processing backend

**Answer:** A

**Explanation:**

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

### NEW QUESTION 101

- (Exam Topic 1)

A company uses AWS Organizations for a multi-account setup in the AWS Cloud. The company uses AWS Control Tower for governance and uses AWS Transit Gateway for VPC connectivity across accounts.

In an AWS application account, the company's application team has deployed a web application that uses AWS Lambda and Amazon RDS. The company's database administrators have a separate DBA account and use the account to centrally manage all the databases across the organization. The database administrators use an Amazon EC2 instance that is deployed in the DBA account to access an RDS database that is deployed in the application account.

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

Which solution will meet these requirements?

- A. Use AWS Resource Access Manager (AWS RAM) to share the secrets from the application account with the DBA account
- B. In the DBA account, create an IAM role that is named DBA-Admi
- C. Grant the role the required permissions to access the shared secret
- D. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.
- E. In the application account, create an IAM role that is named DBA-Secre
- F. Grant the role the required permissions to access the secret
- G. In the DBA account, create an IAM role that is named DBA-Admi
- H. Grant the DBA-Admin role the required permissions to assume the DBA-Secret role in the application account
- I. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.
- J. In the DBA account, create an IAM role that is named DBA-Admi
- K. Grant the role the required permissions to access the secrets and the default AWS managed key in the application account
- L. In the application account, attach resource-based policies to the key to allow access from the DBA account
- M. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.
- N. In the DBA account, create an IAM role that is named DBA-Admi
- O. Grant the role the required permissions to access the secrets in the application account
- P. Attach an SCP to the application account to allow access to the secrets from the DBA account
- Q. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.

**Answer:** B

**Explanation:**

➤ Option B is correct because creating an IAM role in the application account that has permissions to access the secrets and creating an IAM role in the DBA account that has permissions to assume the role in the application account eliminates the need to manually share the secrets. This approach uses cross-account IAM roles to grant access to the secrets in the application account. The database administrators can assume the role in the application account from their EC2 instance in the DBA

account and retrieve the secrets without having to store them locally or share them manually<sup>2</sup>

References: 1: <https://docs.aws.amazon.com/ram/latest/userguide/what-is.html> 2:

[https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial\\_cross-account-with-roles.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html) 3:

<https://docs.aws.amazon.com/kms/latest/developerguide/concepts.html> : [https://docs.aws.amazon.com/secretsmanager/latest/userguide/tutorials\\_basic.html](https://docs.aws.amazon.com/secretsmanager/latest/userguide/tutorials_basic.html) :

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

### NEW QUESTION 102

- (Exam Topic 1)

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

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

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

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

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

**Answer:** D

**Explanation:**

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

An SCP (Service Control Policy) is a policy that you can use to set fine-grained permissions for your AWS

accounts within your organization. You can use SCPs to set permissions for the root user of an account and to delegate permissions to IAM users and roles in the accounts. You can use SCPs to set permissions that allow or deny access to specific services, actions, and resources.

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

the source will be denied, thus meeting the requirement.

Reference: [https://docs.aws.amazon.com/organizations/latest/userguide/orgs\\_manage\\_policies\\_scp.html](https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp.html)

[https://docs.aws.amazon.com/IAM/latest/UserGuide/access\\_policies\\_condition-keys.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_condition-keys.html)

#### NEW QUESTION 105

- (Exam Topic 1)

A company recently deployed an application on AWS. The application uses Amazon DynamoDB. The company measured the application load and configured the RCUs and WCUs on the DynamoDB table to

match the expected peak load. The peak load occurs once a week for a 4-hour period and is double the average load. The application load is close to the average load for the rest of the week. The access pattern includes many more writes to the table than reads of the table.

A solutions architect needs to implement a solution to minimize the cost of the table. Which solution will meet these requirements?

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

**Answer: D**

#### Explanation:

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

#### NEW QUESTION 110

- (Exam Topic 1)

A company is running applications on AWS in a multi-account environment. The company's sales team and marketing team use separate AWS accounts in AWS Organizations.

The sales team stores petabytes of data in an Amazon S3 bucket. The marketing team uses Amazon QuickSight for data visualizations. The marketing team needs access to data that the sales team stores in the S3 bucket. The company has encrypted the S3 bucket with an AWS Key Management Service (AWS KMS) key. The marketing team has already created the IAM service role for QuickSight to provide QuickSight access in the marketing AWS account. The company needs a solution that will provide secure access to the data in the S3 bucket across AWS accounts.

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

- A. Create a new S3 bucket in the marketing account
- B. Create an S3 replication rule in the sales account to copy the objects to the new S3 bucket in the marketing account
- C. Update the QuickSight permissions in the marketing account to grant access to the new S3 bucket.
- D. Create an SCP to grant access to the S3 bucket to the marketing account
- E. Use AWS Resource Access Manager (AWS RAM) to share the KMS key from the sales account with the marketing account
- F. Update the QuickSight permissions in the marketing account to grant access to the S3 bucket.
- G. Update the S3 bucket policy in the marketing account to grant access to the QuickSight role
- H. Create a KMS grant for the encryption key that is used in the S3 bucket
- I. Grant decrypt access to the QuickSight role
- J. Update the QuickSight permissions in the marketing account to grant access to the S3 bucket.
- K. Create an IAM role in the sales account and grant access to the S3 bucket
- L. From the marketing account, assume the IAM role in the sales account to access the S3 bucket
- M. Update the QuickSight role, to create a trust relationship with the new IAM role in the sales account.

**Answer: D**

#### Explanation:

Create an IAM role in the sales account and grant access to the S3 bucket. From the marketing account, assume the IAM role in the sales account to access the S3 bucket. Update the QuickSight role, to create a trust relationship with the new IAM role in the sales account.

This approach is the most secure way to grant cross-account access to the data in the S3 bucket while minimizing operational overhead. By creating an IAM role in the sales account, the marketing team can assume the role in their own account, and have access to the S3 bucket. And updating the QuickSight role, to create a trust relationship with the new IAM role in the sales account will grant the marketing team to access the data in the S3 bucket and use it for data visualization using QuickSight.

AWS Resource Access Manager (AWS RAM) also allows sharing of resources between accounts, but it would require additional management and configuration to set up the sharing, which would increase operational overhead.

Using S3 replication would also replicate the data to the marketing account, but it would not provide the marketing team access to the original data, and also it would increase operational overhead with managing the replication process.

IAM roles and policies, KMS grants and trust relationships are a powerful combination for managing cross-account access in a secure and efficient manner. References:

- > [AWS IAM Roles](#)
- > [AWS KMS - Key Grants](#)
- > [AWS RAM](#)

#### NEW QUESTION 115

.....

## Thank You for Trying Our Product

### We offer two products:

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

2nd - Questions and Answers in PDF Format

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

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

**100% Actual & Verified — Instant Download, Please Click**  
[Order The AWS-Certified-Solutions-Architect-Professional Practice Test Here](#)