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 1AM 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
B. In Account
C. set the S3 bucket policy to the following:

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

D. In Account
E. 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
G. set the permissions of User_DataProcessor to the following:

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

H. In Account Bt 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 I month, the company reviews its EFS storage costs and notices an increase in the EFS backup costs. The EFS backup cold storage produces almost double the cost of the EFS warm backup storage.
What should the solutions architect do to optimize the cost?

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

Answer: A

Explanation:
The cost of EFS backup cold storage is $0.01 per GB-month, whereas the cost of EFS backup warm storage is $0.05 per GB-month1. Therefore, moving the backups to cold storage as soon as possible will reduce the storage cost. However, cold storage backups must be retained for a minimum of 90 day2s, otherwise they incur a pro-rated charge equal to the storage charge for the remaining days1. Therefore, setting the backup retention period to 30 days will incur a penalty of 60 days of cold storage cost for each backup deleted. This penalty will still be lower than keeping the backups in warm storage for 7 days and then in cold storage for 83 days, which is the current configuration. Therefore, option A is the most cost-effective solution.

NEW QUESTION 3
- (Exam Topic 2)
A company is running a two-tier web-based application in an on-premises data center. The application layer consists of a single server running a stateful application. The application connects to a PostgreSQL database running on a separate server. The application's user base is expected to grow significantly, so the company is migrating the application and database to AWS. The solution will use Amazon Aurora PostgreSQL, Amazon EC2 Auto Scaling, and Elastic Load Balancing.
Which solution will provide a consistent user experience that will allow the application and database tiers to scale?

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

Answer: C

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

NEW QUESTION 4
- (Exam Topic 2)
A company needs to optimize the cost of an AWS environment that contains multiple accounts in an organization in AWS Organizations. The company conducted cost optimization activities 3 years ago and purchased Amazon EC2 Standard Reserved Instances that recently expired.
The company needs EC2 instances for 3 more years. Additionally, the company has deployed a new serverless workload.
Which strategy will provide the company with the MOST cost savings?

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

Answer: A

Explanation:
The company should purchase the same Reserved Instances for an additional 3-year term with All Upfront payment. The company should purchase a 3-year Compute Savings Plan with All Upfront payment in the management account to cover any additional compute costs. This solution will provide the company with the most cost savings because Reserved Instances and Savings Plans are both pricing models that offer significant discounts compared to On-Demand pricing. Reserved Instances are commitments to use a specific instance type and size in a single Region for a one- or three-year term. You can choose between three payment options:
No Upfront, Partial Upfront, or All Upfront. The more you pay upfront, the greater the discoun1t. Savings Plans are flexible pricing models that offer low prices on EC2 instances, Fargate, and Lambda usage, in exchange for a commitment to a consistent amount of usage (measured in $/hour) for a one- or three-year term. You can choose between two types of Savings Plans: Compute Savings Plans and EC2 Instance Savings Plans. Compute Savings Plans apply to any EC2 instance regardless of Region, instance family, operating system, or tenancy, including those that are part of EMR, ECS, or EKS clusters, or launched by Fargate or Lambda. EC2 Instance Savings Plans apply to a specific instance family within a Region and provide the most savings2. By purchasing the same Reserved Instances for an additional 3-year term with All Upfront payment, the company can lock in the lowest possible price for its EC2 instances that run continuously for 3 years. By purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account, the company can benefit from additional discounts on any other compute usage across its member accounts.
The other options are not correct because:

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

instance family within a Region and does not cover Fargate or Lambda usage. Also, purchasing a Savings Plan in each member account would not allow the company to share the benefits of unused Savings Plan discounts across its organization.
References:

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

NEW QUESTION 5
- (Exam Topic 2)
A company has loT sensors that monitor traffic patterns throughout a large city. The company wants to read and collect data from the sensors and perform aggregations on the data.
A solutions architect designs a solution in which the loT devices are streaming to Amazon Kinesis Data Streams. Several applications are reading from the stream. However, several consumers are experiencing throttling and are periodically and are periodically encountering a RealProvisioned Throughput Exceeded error.
Which actions should the solution architect take to resolve this issue? (Select THREE.)

A. Reshard the stream to increase the number of shards s in the stream.
B. Use the Kinesis Producer Library KPL). Adjust the polling frequency.
C. Use consumers with the enhanced fan-out feature.
D. Reshard the stream to reduce the number of shards in the stream.
E. Use an error retry and exponential backoff mechanism in the consumer logic.
F. Configure the stream to use dynamic partitioning.

Answer: ACE

Explanation:
https://repost.aws/knowledge-center/kinesis-readprovisionedthroughputexceeded Follow Data Streams best practices
To mitigate ReadProvisionedThroughputExceeded exceptions, apply these best practices:
• Reshard your stream to increase the number of shards in the stream.
• Use consumers with enhanced fan-out. For more information about enhanced fan-out, see Developing custom consumers with dedicated throughput (enhanced fan-out).
• Use an error retry and exponential backoff mechanism in the consumer logic if ReadProvisionedThroughputExceeded exceptions are encountered. For consumer applications that use an AWS SDK, the requests are retried by default.

NEW QUESTION 6
- (Exam Topic 2)
A company wants to use AWS for disaster recovery for an on-premises application. The company has hundreds of Windows-based servers that run the application. All the servers mount a common share.
The company has an RTO of 15 minutes and an RPO of 5 minutes. The solution must support native failover and fallback capabilities.
Which solution will meet these requirements MOST cost-effectively?

A. Create an AWS Storage Gateway File Gatewa
B. Schedule daily Windows server backup
C. Save the data lo Amazon S3. During a disaster, recover the on-premises servers from the backu
D. During failbac
E. run the on-premises servers on Amazon EC2 instances.
F. Create a set of AWS CloudFormation templates to create infrastructur
G. Replicate all data to Amazon Elastic File System (Amazon EFS) by using AWS DataSyn
H. During a disaster, use AWS CodePipeline to deploy the templates to restore the on-premises server
I. Fail back the data by using DataSync.
J. Create an AWS Cloud Development Kit (AWS CDK) pipeline to stand up a multi-site active-active environment on AW
K. Replicate data into Amazon S3 by using the s3 sync comman
L. During a disaster, swap DNS endpoints to point to AW
M. Fail back the data by using the s3 sync command.
N. Use AWS Elastic Disaster Recovery to replicate the on-premises server
O. Replicate data to an Amazon FSx for Windows File Server file system by using AWS DataSyn
P. Mount the file system to AWS server
Q. During a disaster, fail over the on-premises servers to AW
R. Fail back to new or existing servers by using Elastic Disaster Recovery.

Answer: D

NEW QUESTION 7
- (Exam Topic 2)
A solutions architect needs to improve an application that is hosted in the AWS Cloud. The application uses an Amazon Aurora MySQL DB instance that is experiencing overloaded connections. Most of the application's operations insert records into the database. The application currently stores credentials in a text-based configuration file.
The solutions architect needs to implement a solution so that the application can handle the current connection load. The solution must keep the credentials secure and must provide the ability to rotate the credentials automatically on a regular basis.
Which solution will meet these requirements?

A. Deploy an Amazon RDS Proxy layer in front of the DB 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 8
- (Exam Topic 2)
A company has an application that runs on Amazon EC2 instances in an Amazon EC2 Auto Scaling group. The company uses AWS CodePipeline to deploy the application. The instances that run in the Auto Scaling group are constantly changing because of scaling events.
When the company deploys new application code versions, the company installs the AWS CodeDeploy agent on any new target EC2 instances and associates the instances with the CodeDeploy deployment group. The application is set to go live within the next 24 hours.
What should a solutions architect recommend to automate the application deployment process with the LEAST amount of operational overhead?

A. Configure Amazon EventBridge to invoke an AWS Lambda function when a new EC2 instance is launched into the Auto Scaling grou
B. Code the Lambda function to associate the EC2 instances with the CodeDeploy deployment group.
C. Write a script to suspend Amazon EC2 Auto Scaling operations before the deployment of new code When the deployment is complete, create a new AMI and configure the Auto Scaling group's launch template to use the new AMI for new launche
D. Resume Amazon EC2 Auto Scaling operations.
E. Create a new AWS CodeBuild project that creates a new AMI that contains the new code Configure CodeBuild to update the Auto Scaling group's launch template to the new AM
F. Run an Amazon EC2 Auto Scaling instance refresh operation.
G. Create a new AMI that has the CodeDeploy agent installe
H. Configure the Auto Scaling group's launch template to use the new AM
I. Associate the CodeDeploy deployment group with the Auto Scaling group instead of the EC2 instances.

Answer: D

Explanation:
https://docs.aws.amazon.com/codedeploy/latest/userguide/integrations-aws-auto-scaling.html

NEW QUESTION 9
- (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 I business day, with loss of no more than I 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 10
- (Exam Topic 2)
A company is designing a new website that hosts static content. The website will give users the ability to upload and download large files. According to company requirements, all data must be encrypted in transit and at rest. A solutions architect is building the solution by using Amazon S3 and Amazon CloudFront.
Which combination of steps will meet the encryption requirements? (Select THREE.)

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

Answer: ACE

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

NEW QUESTION 10
- (Exam Topic 2)
A company 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 theRemoveAccountFromOrganization operation in the Organizations API.
D. Sign in to the new developer organization's management account and create a placeholder member account that acts as a target for the developer account migration.
E. Call the InviteAccountToOrganization operation in the Organizations API from the new developer organization's management account to send invitations to the developer accounts.
F. Have each developer sign in to their account and confirm to join the new developer organization.

Answer: BEF

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

NEW QUESTION 12
- (Exam Topic 2)
A company has developed a hybrid solution between its data center and AWS. The company uses Amazon VPC and Amazon EC2 instances that send application togs to Amazon CloudWatch. The EC2 instances read data from multiple relational databases that are hosted on premises.
The company wants to monitor which EC2 instances are connected to the databases in near-real time. The company already has a monitoring solution that uses Splunk on premises. A solutions architect needs to determine how to send networking traffic to Splunk.
How should the solutions architect meet these requirements?

A. Enable VPC flows logs, and send them to CloudWatc
B. Create an AWS Lambda function to periodically export the CloudWatch logs to an Amazon S3 bucket by using the pre-defined export functio
C. Generate ACCESS_KEY and SECRET_KEY AWS credential
D. Configure Splunk to pull the logs from the S3 bucket by using those credentials.
E. Create an Amazon Kinesis Data Firehose delivery stream with Splunk as the destinatio
F. Configure a pre-processing AWS Lambda function with a Kinesis Data Firehose stream processor that extracts individual log events from records sent by CloudWatch Logs subscription filter
G. Enable VPC flows logs, and send them to CloudWatc
H. Create a CloudWatch Logs subscription that sends log events to the Kinesis Data Firehose delivery stream.
I. Ask the company to log every request that is made to the databases along with the EC2 instance IP addres
J. Export the CloudWatch logs to an Amazon S3 bucke
K. Use Amazon Athena to query the logs grouped by database nam
L. Export Athena results to another S3 bucke
M. Invoke an AWS Lambdafunction to automatically send any new file that is put in the S3 bucket to Splunk.
N. Send the CloudWatch logs to an Amazon Kinesis data stream with Amazon Kinesis Data Analytics for SOL Application
O. Configure a 1 -minute sliding window to collect the event
P. Create a SQL query that uses the anomaly detection template to monitor any networking traffic anomalies in near-real tim
Q. Send the result to an Amazon Kinesis Data Firehose delivery stream with Splunk as the destination.

Answer: B

Explanation:
https://docs.aws.amazon.com/firehose/latest/dev/creating-the-stream-to-splunk.html

NEW QUESTION 14
- (Exam Topic 2)
A company is migrating a document processing workload to AWS. The company has updated many applications to natively use the Amazon S3 API to store, retrieve, and modify documents that a processing server generates at a rate of approximately 5 documents every second. After the document processing is finished, customers can download the documents directly from Amazon S3.
During the migration, the company discovered that it could not immediately update the processing server that generates many documents to support the S3 API. The server runs on Linux and requires fast local access to the files that the server generates and modifies. When the server finishes processing, the files must be available to the public for download within 30 minutes.
Which solution will meet these requirements with the LEAST amount of effort?

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

Answer: C

Explanation:
The company should configure Amazon FSx for Lustre with an import and export policy. The company should link the new file system to an S3 bucket. The company should install the Lustre client and mount the document store to an Amazon EC2 instance by using NFS. This solution will meet the requirements with

the least amount of effort because Amazon FSx for Lustre is a fully managed service that provides a high-performance file system optimized for fast processing of workloads such as machine learning, high performance computing, video processing, financial modeling, and electronic design automation1. Amazon FSx for Lustre can be linked to an S3 bucket and can import data from and export data to the bucket2. The import and export policy can be configured to automatically import new or changed objects from S3 and export new or changed files to S33. This will ensure that the files are available to the public for download within 30 minutes. Amazon FSx for Lustre supports NFS version 3.0 protocol for Linux clients.
The other options are not correct because:

- Migrating the application to an AWS Lambda function would require a lot of effort and may not be feasible for the existing server that generates many documents. Lambda functions have limitations on execution time, memory, disk space, and network bandwidth.
- Setting up an Amazon S3 File Gateway would not work because S3 File Gateway does not support write-back caching, which means that files written to the file share are uploaded to S3 immediately and are not available locally until they are downloaded again. This would not provide fast local access to the files that the server generates and modifies.
- Configuring AWS DataSync to connect to an Amazon EC2 instance would not meet the requirement of making the files available to the public for download within 30 minutes. DataSync is a service that transfers data between on-premises storage systems and AWS storage services over the internet or AWS Direct Connect. DataSync tasks can be scheduled to run at specific times or intervals, but they are not triggered by file changes.

References:

- https://aws.amazon.com/fsx/lustre/
- https://docs.aws.amazon.com/fsx/latest/LustreGuide/create-fs-linked-data-repo.html
- https://docs.aws.amazon.com/fsx/latest/LustreGuide/import-export-data-repositories.html
- https://docs.aws.amazon.com/fsx/latest/LustreGuide/mounting-on-premises.html
- https://docs.aws.amazon.com/lambda/latest/dg/gettingstarted-limits.html
- https://docs.aws.amazon.com/storagegateway/latest/userguide/StorageGatewayConcepts.html
- https://docs.aws.amazon.com/datasync/latest/userguide/what-is-datasync.html

NEW QUESTION 19
- (Exam Topic 2)
A company's interactive web application uses an Amazon CloudFront distribution to serve images from an Amazon S3 bucket. Occasionally, third-party tools ingest corrupted images into the S3 bucket. This image corruption causes a poor user experience in the application later. The company has successfully implemented and tested Python logic to detect corrupt images.
A solutions architect must recommend a solution to integrate the detection logic with minimal latency between the ingestion and serving.
Which solution will meet these requirements?

A. Use a Lambda@Edge function that is invoked by a viewer-response event.
B. Use a Lambda@Edge function that is invoked by an origin-response event.
C. Use an S3 event notification that invokes an AWS Lambda function.
D. Use an S3 event notification that invokes an AWS Step Functions state machine.

Answer: B

Explanation:
This solution will allow the detection logic to be run as soon as the image is uploaded to the S3 bucket, before it is served to users via the CloudFront distribution. This way, the detection logic can quickly identify any corrupted images and prevent them from being served to users, minimizing latency between ingestion and serving.
Reference: AWS Lambda@Edge documentation:
https://docs.aws.amazon.com/lambda/latest/dg/lambda-edge.html You can use Lambda@Edge to run your code in response to CloudFront events, such as a viewer request, an origin request, a response, or an error.

NEW QUESTION 24
- (Exam Topic 2)
A company uses AWS Organizations with a single OU named Production to manage multiple accounts All accounts are members of the Production OU Administrators use deny list SCPs in the root of the organization to manage access to restricted services.
The company recently acquired a new business unit and invited the new unit's existing AWS account to the organization Once onboarded the administrators of the new business unit discovered that they are not able to update existing AWS Config rules to meet the company's policies.
Which option will allow administrators to make changes and continue to enforce the current policies without introducing additional long-term maintenance?

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

Answer: D

Explanation:
An SCP at a lower level can't add a permission after it is blocked by an SCP at a higher level. SCPs can only filter; they never add permissions. SO you need to create a new OU for the new account assign an SCP, and move the root SCP to Production OU. Then move the new account to production OU when AWS config is done.

NEW QUESTION 29
- (Exam Topic 2)
A company is running an application that uses an Amazon ElastiCache for Redis cluster as a caching layer A recent security audit revealed that the company has configured encryption at rest for ElastiCache However the company did not configure ElastiCache to use encryption in transit Additionally, users can access the

cache without authentication
A solutions architect must make changes to require user authentication and to ensure that the company is using end-to-end encryption
Which solution will meet these requirements?

A. Create an AUTH token Store the token in AWS System Manager Parameter Store, as an encrypted parameter Create a new cluster with AUTH and configure encryption in transit Update the application toretrieve the AUTH token from Parameter Store when necessary and to use the AUTH token for authentication
B. Create an AUTH token Store the token in AWS Secrets Manager Configure the existing cluster to use the AUTH token and configure encryption in transit Update the application to retrieve the AUTH token from Secrets Manager when necessary and to use the AUTH token for authentication.
C. Create an SSL certificate Store the certificate in AWS Secrets Manager Create a new cluster and configure encryption in transit Update the application to retrieve the SSL certificate from Secrets Manager when necessary and to use the certificate for authentication.
D. Create an SSL certificate Store the certificate in AWS Systems Manager Parameter Store, as an encrypted advanced parameter Update the existing cluster to configure encryption in transit Update the application to retrieve the SSL certificate from Parameter Store when necessary and to use the certificate for authentication

Answer: B

Explanation:
Creating an AUTH token and storing it in AWS Secrets Manager and configuring the existing cluster to use the AUTH token and configure encryption in transit, and updating the application to retrieve the AUTH token from Secrets Manager when necessary and to use the AUTH token for authentication, would meet the requirements for user authentication and end-to-end encryption.
AWS Secrets Manager is a service that enables you to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. Secrets Manager also enables you to encrypt the data and ensure that only authorized users and applications can access it.
By configuring the existing cluster to use the AUTH token and encryption in transit, all data will be encrypted as it is sent over the network, providing additional security for the data stored in ElastiCache.
Additionally, by updating the application to retrieve the AUTH token from Secrets Manager when necessary and to use the AUTH token for authentication, it ensures that only authorized users and applications can access the cache.
Reference:
AWS Secrets Manager documentation: https://aws.amazon.com/secrets-manager/ Encryption in transit for ElastiCache: https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/encryption.html
Authentication and Authorization for ElastiCache: https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/accessing-elasticache.html

NEW QUESTION 34
- (Exam Topic 2)
A company needs to migrate its customer transactions database from on premises to AWS. The database resides on an Oracle DB instance that runs on a Linux server. According to a new security requirement, the company must rotate the database password each year.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: B

NEW QUESTION 37
- (Exam Topic 2)
A company has a few AWS accounts for development and wants to move its production application to AWS. The company needs to enforce Amazon Elastic Block Store (Amazon EBS) encryption at rest current production accounts and future production accounts only. The company needs a solution that includes built-in blueprints and guardrails.
Which combination of steps will meet these requirements? (Choose three.)

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

Answer: CDF

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

NEW QUESTION 42
- (Exam Topic 2)
A software-as-a-service (SaaS) provider exposes APIs through an Application Load Balancer (ALB). The ALB connects to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster that is deployed in the

us-east-I Region. The exposed APIs contain usage of a few non-standard REST methods: LINK, UNLINK, LOCK, and UNLOCK.
Users outside the United States are reporting long and inconsistent response times for these APIs. A solutions architect needs to resolve this problem with a solution that minimizes operational overhead.
Which solution meets these requirements?

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

Answer: C

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

NEW QUESTION 43
- (Exam Topic 2)
A solutions architect is planning to migrate critical Microsoft SOL Server databases to AWS. Because the databases are legacy systems, the solutions architect will move the databases to a modern data architecture. The solutions architect must migrate the databases with near-zero downtime.
Which solution will meet these requirements?

A. Use AWS Application Migration Service and the AWS Schema Conversion Tool (AWS SCT). Perform an In-place upgrade before the migratio
B. Export the migrated data to Amazon Aurora Serverless after cutove
C. Repoint the applications to Amazon Aurora.
D. Use AWS Database Migration Service (AWS DMS) to Rehost the databas
E. Set Amazon S3 as a target.Set up change data capture (CDC) replicatio
F. When the source and destination are fully synchronized, load the data from Amazon S3 into an Amazon RDS for Microsoft SQL Server DB Instance.
G. Use native database high availability tools Connect the source system to an Amazon RDS for Microsoft SQL Server DB instance Configure replication accordingl
H. When data replication is finished, transition the workload to an Amazon RDS for Microsoft SQL Server DB instance.
I. Use AWS Application Migration Servic
J. Rehost the database server on Amazon EC2. When data replication is finished, detach the database and move the database to an Amazon RDS for Microsoft SQL Server DB instanc
K. Reattach the database and then cut over all networking.

Answer: B

Explanation:
AWS DMS can migrate data from a source database to a target database in AWS, using change data capture (CDC) to replicate ongoing changes and keep the databases in sync. Setting Amazon S3 as a target allows storing the migrated data in a durable and cost-effective storage service. When the source and destination are fully synchronized, the data can be loaded from Amazon S3 into an Amazon RDS for Microsoft SQL Server DB instance, which is a managed database service that simplifies database administration tasks. References:

- https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Source.SQLServer.html
- https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.S3.html
- https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_SQLServer.html

NEW QUESTION 47
- (Exam Topic 2)
A solutions architect needs to review the design of an Amazon EMR cluster that is using the EMR File System (EMRFS). The cluster performs tasks that are critical to business needs. The cluster is running Amazon EC2 On-Demand Instances at all times tor all task, primary, and core nodes. The EMR tasks run each morning, starting at 1 ;00 AM. and take 6 hours to finish running. The amount of time to complete the processing is not a priority because the data is not referenced until late in the day.
The solutions architect must review the architecture and suggest a solution to minimize the compute costs. Which solution should the solutions architect recommend to meet these requirements?

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

Answer: A

Explanation:
Amazon EC2 Spot Instances offer spare compute capacity at steep discounts compared to On-Demand prices. Spot Instances can be interrupted by EC2 with two minutes of notification when EC2 needs the capacity back. Amazon EMR can handle Spot interruptions gracefully by decommissioning the nodes and

redistributing the tasks to other nodes. By launching all nodes on Spot Instances in an instance fleet, the solutions architect can minimize the compute costs of the EMR cluster. An instance fleet is a collection of EC2 instances with different types and sizes that EMR automatically provisions to meet a defined target capacity. By terminating the cluster when the processing is completed, the solutions architect can avoid paying for idle resources. References:

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

NEW QUESTION 52
- (Exam Topic 2)
A company is running an application on Amazon EC2 instances in the AWS Cloud. The application is using a MongoDB database with a replica set as its data tier. The MongoDB database is installed on systems in the company's on-premises data center and is accessible through an AWS Direct Connect connection to the data center environment.
A solutions architect must migrate the on-premises MongoDB database to Amazon DocumentDB (with MongoDB compatibility).
Which strategy should the solutions architect choose to perform this migration?

A. Create a fleet of EC2 instance
B. Install MongoDB Community Edition on the EC2 instances, and create a databas
C. Configure continuous synchronous replication with the database that is running in theon-premises data center.
D. Create an AWS Database Migration Service (AWS DMS) replication instanc
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 databas
F. Create and run a DMS migration task.
G. Create a data migration pipeline by using AWS Data Pipelin
H. Define data nodes for the on-premises MongoDB database and the Amazon DocumentDB databas
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 54
- (Exam Topic 2)
A company hosts a blog post application on AWS using Amazon API Gateway, Amazon DynamoDB, and AWS Lambda. The application currently does not use API keys to authorize requests. The API model is as follows: GET/posts/[postid] to get post details GET/users[userid] to get user details GET/comments/[commentid] to get comments details
The company has noticed users are actively discussing topics in the comments section, and the company wants to increase user engagement by marking the comments appears in real time.
Which design should be used to reduce comment latency and improve user experience?

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

Answer: C

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

NEW QUESTION 56
- (Exam Topic 2)
A company uses an AWS CodeCommit repository The company must store a backup copy of the data that is in the repository in a second AWS Region
Which solution will meet these requirements?

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

Answer: B

Explanation:
AWS Backup is a fully managed service that makes it easy to centralize and automate the creation, retention, and restoration of backups across AWS services. It provides a way to schedule automatic backups for CodeCommit repositories on an hourly basis. Additionally, it also supports cross-Region replication, which allows you to copy the backups to a second Region for disaster recovery.
By using AWS Backup, the company can set up an automatic and regular backup schedule for the CodeCommit repository, ensuring that the data is regularly backed up and stored in a second Region. This can provide a way to recover quickly from any disaster event that might occur.
Reference:
AWS Backup documentation: https://aws.amazon.com/backup/ AWS Backup for AWS CodeCommit documentation: https://aws.amazon.com/about-aws/whats-new/2020/07/aws-backup-now-supports-aws-codecommit-repositorie

NEW QUESTION 60
- (Exam Topic 2)
A company has an on-premises Microsoft SOL 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 bucke
B. Deploy an AWS Storage Gateway file gateway within the VPC that Is connected to the Direct Connect connectio
C. Create a new SMB file shar
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 connectio
F. Create a new SMB file shar
G. Write nightly database exports to an SMB file share on the Amazon FSx file syste
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 connectio
J. Create a new SMB file shar
K. Write nightly database exports to an SMB file share on the Amazon FSx file syste
L. Enable nightly backups.
M. Create a new S3 bucke
N. Deploy an AWS Storage Gateway volume gateway within the VPC that Is connected to the Direct Connect connectio
O. Create a new SMB file shar
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 64
- (Exam Topic 2)
A company's public API runs as tasks on Amazon Elastic Container Service (Amazon ECS). The tasks run on AWS Fargate behind an Application Load Balancer (ALB) and are configured with Service Auto Scaling for the tasks based on CPU utilization. This service has been running well for several months.
Recently, API performance slowed down and made the application unusable. The company discovered that a significant number of SQL injection attacks had occurred against the API and that the API service had scaled to its maximum amount.
A solutions architect needs to implement a solution that prevents SQL injection attacks from reaching the ECS API service. The solution must allow legitimate traffic through and must maximize operational efficiency. Which solution meets these requirements?

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

Answer: C

Explanation:
The company should create a new AWS WAF web ACL. The company should add a new rule that blocks requests that match the SQL database rule group. The company should set the web ACL to allow all other traffic that does not match those rules. The company should attach the web ACL to the ALB in front of the ECS tasks. This solution will meet the requirements because AWS WAF is a web application firewall that lets you monitor and control web requests that are forwarded to your web applications. You can use AWS WAF to define customizable web security rules that control which traffic can access your web applications and which traffic should be blocked1. By creating a new AWS WAF web ACL, the company can create a collection of rules that define the conditions for allowing or blocking web requests. By adding a new rule that blocks requests that match the SQL database rule group, the company can prevent SQL injection attacks from reaching the ECS API service. The SQL database rule group is a managed rule group provided by AWS that contains rules to protect against common SQL injection attack patterns2. By setting the web ACL to allow all other traffic that does not match those rules, the company can ensure that legitimate traffic can access the API service. By attaching the web ACL to the ALB in front of the ECS tasks, the company can apply the web security rules to all requests that are forwarded by the load balancer.
The other options are not correct because:
- Creating a new AWS WAF Bot Control implementation would not prevent SQL injection attacks from reaching the ECS API service. AWS WAF Bot Control is a feature that gives you visibility and control over common and pervasive bot traffic that can consume excess resources, skew metrics, cause downtime, or perform other undesired activities. However, it does not protect against SQL injection attacks, which are malicious attempts to execute unauthorized SQL statements against your database3.
- Creating a new AWS WAF web ACL to monitor the HTTP requests and HTTPS requests that are forwarded to the ALB in front of the ECS tasks would not prevent SQL injection attacks from reaching
the ECS API service. Monitoring mode is a feature that enables you to evaluate how your rules would perform without actually blocking any requests. However, this mode does not provide any protection against attacks, as it only logs and counts requests that match your rules4.
- Creating a new AWS WAF web ACL and creating a new empty IP set in AWS WAF would not prevent SQL injection attacks from reaching the ECS API service. An IP set is a feature that enables you to specify a list of IP addresses or CIDR blocks that you want to allow or block based on their source IP address. However, this approach would not be effective or efficient against SQL injection attacks, as it would require constantly updating the IP set with new IP addresses of attackers, and it would not block attackers who use proxies or VPNs.
References:
- https://aws.amazon.com/waf/
- https://docs.aws.amazon.com/waf/latest/developerguide/waf-bot-control.html

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

NEW QUESTION 68
- (Exam Topic 2)
A company has a critical application in which the data tier is deployed in a single AWS Region. The data tier uses an Amazon DynamoDB table and an Amazon Aurora MySQL DB cluster. The current Aurora MySQL engine version supports a global database. The application tier is already deployed in two Regions. Company policy states that critical applications must have application tier components and data tier components deployed across two Regions. The RTO and RPO must be no more than a few minutes each. A solutions architect must recommend a solution to make the data tier compliant with company policy.
Which combination of steps will meet these requirements? (Choose two.)

A. Add another Region to the Aurora MySQL DB cluster
B. Add another Region to each table in the Aurora MySQL DB cluster
C. Set up scheduled cross-Region backups for the DynamoDB table and the Aurora MySQL DB cluster
D. Convert the existing DynamoDB table to a global table by adding another Region to its configuration
E. Use Amazon Route 53 Application Recovery Controller to automate database backup and recovery to the secondary Region

Answer: AD

Explanation:
The company should use Amazon Aurora global database and Amazon DynamoDB global table to deploy the data tier components across two Regions. Amazon Aurora global database is a feature that allows a single
Aurora database to span multiple AWS Regions, enabling low-latency global reads and fast recovery from Region-wide outages1. Amazon DynamoDB global table is a feature that allows a single DynamoDB table to span multiple AWS Regions, enabling low-latency global reads and writes and fast recovery from Region-wide outages2.
References:
- https://aws.amazon.com/rds/aurora/global-database/
- https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/globaltables_HowItWorks.html
- https://aws.amazon.com/route53/application-recovery-controller/

NEW QUESTION 70
- (Exam Topic 2)
A company runs an application in an on-premises data center. The application gives users the ability to upload media files. The files persist in a file server. The web application has many users. The application server is overutilized, which causes data uploads to fail occasionally. The company frequently adds new storage to the file server. The company wants to resolve these challenges by migrating the application to AWS.
Users from across the United States and Canada access the application. Only authenticated users should have the ability to access the application to upload files. The company will consider a solution that refactors the application, and the company needs to accelerate application development.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use AWS Application Migration Service to migrate the application server to Amazon EC2 instances.Create an Auto Scaling group for the EC2 instance
B. Use an Application Load Balancer to distribute the request
C. Modify the application to use Amazon S3 to persist the file
D. Use Amazon Cognito to authenticate users.
E. Use AWS Application Migration Service to migrate the application server to Amazon EC2 instances.Create an Auto Scaling group for the EC2 instance
F. Use an Application Load Balancer to distribute the request
G. Set up AWS IAM Identity Center (AWS Single Sign-On) to give users the ability to sign in to the applicatio
H. Modify the application to use Amazon S3 to persist the files.
I. Create a static website for uploads of media file
J. Store the static assets in Amazon S3. Use AWS AppSync to create an AP
K. Use AWS Lambda resolvers to upload the media files to Amazon S3. Use Amazon Cognito to authenticate users.
L. Use AWS Amplify to create a static website for uploads of media file
M. Use Amplify Hosting to serve the website through Amazon CloudFron
N. Use Amazon S3 to store the uploaded media file
O. Use Amazon Cognito to authenticate users.

Answer: D

Explanation:
The company should use AWS Amplify to create a static website for uploads of media files. The company should use Amplify Hosting to serve the website through Amazon CloudFront. The company should use Amazon S3 to store the uploaded media files. The company should use Amazon Cognito to authenticate users.
This solution will meet the requirements with the least operational overhead because AWS Amplify is a complete solution that lets frontend web and mobile developers easily build, ship, and host full-stack applications on AWS, with the flexibility to leverage the breadth of AWS services as use cases evolve. No cloud expertise needed1. By using AWS Amplify, the company can refactor the application to a serverless architecture that reduces operational complexity and costs.
AWS Amplify offers the following features and benefits:
- Amplify Studio: A visual interface that enables you to build and deploy a full-stack app quickly, including frontend UI and backend.
- Amplify CLI: A local toolchain that enables you to configure and manage an app backend with just a few commands.
- Amplify Libraries: Open-source client libraries that enable you to build cloud-powered mobile and web apps.
- Amplify UI Components: Open-source design system with cloud-connected components for building feature-rich apps fast.
- Amplify Hosting: Fully managed CI/CD and hosting for fast, secure, and reliable static and server-side rendered apps.

By using AWS Amplify to create a static website for uploads of media files, the company can leverage Amplify Studio to visually build a pixel-perfect UI and connect it to a cloud backend in clicks. By using Amplify Hosting to serve the website through Amazon CloudFront, the company can easily deploy its web app or website to the fast, secure, and reliable AWS content delivery network (CDN), with hundreds of points of presence globally. By using Amazon S3 to store the uploaded media files, the company can benefit from a highly scalable, durable, and cost-effective object storage service that can handle any amount of data2. By using Amazon Cognito to authenticate users, the company can add user sign-up, sign-in, and access control to its web app with a fully managed service that scales to support millions of users3.
The other options are not correct because:
- Using AWS Application Migration Service to migrate the application server to Amazon EC2 instances would not refactor the application or accelerate

development. AWS Application Migration Service (AWS MGN) is a service that enables you to migrate physical servers, virtual machines (VMs), or cloud servers from any source infrastructure to AWS without requiring agents or specialized tools. However, this would not address the challenges of overutilization and data uploads failures. It would also not reduce operational overhead or costs compared to a serverless architecture.

- Creating a static website for uploads of media files and using AWS AppSync to create an API would not be as simple or fast as using AWS Amplify. AWS AppSync is a service that enables you to create flexible APIs for securely accessing, manipulating, and combining data from one or more data sources. However, this would require more configuration and management than using Amplify Studio and Amplify Hosting. It would also not provide authentication features like Amazon Cognito.
- Setting up AWS IAM Identity Center (AWS Single Sign-On) to give users the ability to sign in to the application would not be as suitable as using Amazon Cognito. AWS Single Sign-On (AWS SSO) is a service that enables you to centrally manage SSO access and user permissions across multiple AWS accounts and business applications. However, this service is designed for enterprise customers who need to manage access for employees or partners across multiple resources. It is not intended for authenticating end users of web or mobile apps.

References:

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

NEW QUESTION 73
- (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 78
- (Exam Topic 2)
A company runs an intranet application on premises. The company wants to configure a cloud backup of the application. The company has selected AWS Elastic Disaster Recovery for this solution.
The company requires that replication traffic does not travel through the public internet. The application also must not be accessible from the internet. The company does not want this solution to consume all available network bandwidth because other applications require bandwidth.
Which combination of steps will meet these requirements? (Select THREE.)

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

Answer: BDE

Explanation:
AWS Elastic Disaster Recovery (AWS DRS) is a service that minimizes downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications using affordable storage, minimal compute, and point-in-time recovery1. Users can set up AWS DRS on their source servers to initiate secure data replication to a staging area subnet in their AWS account, in the AWS Region they select. Users can then launch recovery instances on AWS within minutes, using the most up-to-date server state or a previous point in time.
To configure a cloud backup of the application with AWS DRS, users need to create a VPC that has at least
two public subnets, a virtual private gateway, and an internet gateway. A VPC is a logically isolated section of the AWS Cloud where users can launch AWS

resources in a virtual network that they define2. A public subnet is a subnet that has a route to an internet gateway3. A virtual private gateway is the VPN concentrator on the Amazon side of the Site-to-Site VPN connection4. An internet gateway is a horizontally scaled, redundant, and highly available VPC component that allows communication between instances in the VPC and the internet. Users need to create at least two public subnets for redundancy and high availability. Users need to create a virtual private gateway and attach it to the VPC to enable VPN connectivity between the on-premises network and the target AWS network. Users need to create an internet gateway and attach it to the VPC to enable internet access for the replication servers.
To ensure that replication traffic does not travel through the public internet, users need to create an AWS Direct Connect connection and a Direct Connect gateway between the on-premises network and the target AWS network. AWS Direct Connect is a service that establishes a dedicated network connection from an on-premises network to one or more VPCs. A Direct Connect gateway is a globally available resource that allows users to connect multiple VPCs across different Regions to their on-premises networks using one or more Direct Connect connections. Users need to create an AWS Direct Connect connection between their on-premises network and an AWS Region. Users need to create a Direct Connect gateway and associate it with their VPC and their Direct Connect connection.
To ensure that the application is not accessible from the internet, users need to select the option to use private IP addresses for data replication during configuration of the replication servers. This option configures the replication servers with private IP addresses only, without assigning any public IP addresses or Elastic IP addresses. This way, the replication servers can only communicate with other resources within the VPC or through VPN connections.
Option A is incorrect because creating a VPC that has at least two private subnets, two NAT gateways, and a virtual private gateway is not necessary or cost-effective. A private subnet is a subnet that does not have a route to an internet gateway3. A NAT gateway is a highly available, managed Network Address Translation (NAT) service that enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating connections with those instances. Users do not need to create private subnets or NAT gateways for this use case, as they can use public subnets with private IP addresses for data replication.
Option C is incorrect because creating an AWS Site-to-Site VPN connection between the on-premises network and the target AWS network will not ensure that replication traffic does not travel through the public
internet. A Site-to-Site VPN connection consists of two VPN tunnels between an on-premises customer
gateway device and a virtual private gateway in your VPC4. The VPN tunnels are encrypted using IPSec protocols, but they still use public IP addresses for communication. Users need to use AWS Direct Connect instead of Site-to-Site VPN for this use case.
Option F is incorrect because selecting the option to ensure that the Recovery instance's private IP address matches the source server's private IP address during configuration of the launch settings for the target servers will not ensure that the application is not accessible from the internet. This option configures the Recovery instance with an identical private IP address as its source server when launched in drills or recovery mode. However, this option does not prevent assigning public IP addresses or Elastic IP addresses to the Recovery instance. Users need to select the option to use private IP addresses for data replication instead.

NEW QUESTION 79
- (Exam Topic 2)
A company is designing an AWS Organizations structure. The company wants to standardize a process to apply tags across the entire organization. The company will require tags with specific values when a user creates a new resource. Each of the company's OUs will have unique tag values.
Which solution will meet these requirements?

A. Use an SCP to deny the creation of resources that do not have the required tag
B. Create a tag policy that Includes the tag values that the company has assigned to each O
C. Attach the tag policies to the OUs.
D. Use an SCP to deny the creation of resources that do not have the required tag
E. Create a tag policy that includes the tag values that the company has assigned to each O
F. Attach the tag policies to the organization's management account.
G. Use an SCP to allow the creation of resources only when the resources have the required tag
H. Create a tag policy that includes the tag values that the company has assigned to each O
I. Attach the tag policies to the OUs.
J. Use an SCP to deny the creation of resources that do not have the required tag
K. Define the list of tags.Attach the SCP to the OUs

Answer: A

Explanation:
https://aws.amazon.com/blogs/mt/implement-aws-resource-tagging-strategy-using-aws-tag-policies-and-service

NEW QUESTION 80
- (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 syste
B. Configure the file system for 75 MiBps of provisioned throughpu
C. Implementreplication to a file system in the DR Region.
D. Deploy a new Amazon FSx for Lustre file syste
E. Configure Bursting Throughput mode for the file syste
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 throughpu
H. Enable Multi-Attach for the EBSvolum
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 thedata from the production file system to the DR file system every 10 minutes.

Answer: A

Explanation:
The company should deploy a new Amazon Elastic File System (Amazon EFS) Multi-AZ file system. The company should configure the file system for 75 MiBps of provisioned throughput. The company should implement replication to a file system in the DR Region. This solution will meet the requirements because Amazon EFS is a serverless, fully elastic file storage service that lets you share file data without provisioning or managing storage capacity and performance. Amazon EFS is built to scale on demand to petabytes without disrupting applications, growing and shrinking automatically as you add and remove files1. By deploying a new Amazon EFS Multi-AZ file system, the company can create a single location for updates to application data for all instances. A Multi-AZ file system replicates data

across multiple Availability Zones (AZs) within a Region, providing high availability and durability2. By configuring the file system for 75 MiBps of provisioned throughput, the company can ensure that it meets the peak operations requirement of 225 MiBps of read throughput. Provisioned throughput is a feature that enables you to specify a level of throughput that the file system can drive independent of the file system's size or burst credit balance3. By implementing replication to a file system in the DR Region, the company can make a copy of the data available in another AWS Region for disaster recovery. Replication is a feature that enables you to replicate data from one EFS file system to another EFS file system across AWS Regions. The replication process has an RPO of less than 1 hour.
The other options are not correct because:

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

References:

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

NEW QUESTION 85
- (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 typ
B. Copy the existing objects to the new S3 bucke
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 typ
E. Use S3 Batch Operations to copy the existing objects to the new S3 bucke
F. Specify SSE-S3.
G. Use AWS CloudHSM to store the encryption key
H. Create a new S3 bucke
I. Use S3 Batch Operations to copy the existing objects to the new S3 bucke
J. Encrypt the objects by using the keys from CloudHSM.
K. Use the S3 Intelligent-Tiering storage class for the S3 bucke
L. Create an S3 Intelligent-Tiering archive configuration to transition objects that are not accessed for 90 days to S3 Glacier Deep Archive.

Answer: B

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

NEW QUESTION 88
- (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 92
- (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 VP
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 AW
D. Provision a transit VIF, and connect it to a Direct Connect gatewa
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 VP
G. Use a transit gateway with dynamic routin
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 Regio
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 96
- (Exam Topic 2)
A solutions architect wants to cost-optimize and appropriately size Amazon EC2 instances in a single AWS account. The solutions architect wants to ensure that the instances are optimized based on CPU, memory, and network metrics.
Which combination of steps should the solutions architect take to meet these requirements? (Choose two.)

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

Answer: BD

Explanation:
AWS Trusted Advisor is a service that provides real-time guidance to help users provision their resources following AWS best practices1. One of the Trusted Advisor checks is “Low Utilization Amazon EC2 Instances”, which identifies EC2 instances that appear to be underutilized based on CPU, network I/O, and disk I/O metrics1. This check can help users optimize the cost and size of their EC2 instances by recommending smaller or more appropriate instance types.
AWS Compute Optimizer is a service that analyzes the configuration and utilization metrics of AWS resources and generates optimization recommendations to reduce the cost and improve the performance of workloads2. Compute Optimizer supports four types of AWS resources: EC2 instances, EBS volumes, ECS services on AWS Fargate, and Lambda functions2. For EC2 instances, Compute Optimizer evaluates the vCPUs, memory, storage, and other specifications, as well as the CPU utilization, network in and out, disk read and write, and other utilization metrics of currently running instances3. It then recommends optimal instance types based on price-performance trade-offs.
Option A is incorrect because purchasing AWS Business Support or AWS Enterprise Support for the account will not directly help with cost-optimization and sizing of EC2 instances. However, these support plans do provide access to more Trusted Advisor checks than the basic support plan1.
Option C is incorrect because installing the Amazon CloudWatch agent and configuring memory metric collection on the EC2 instances will not provide any optimization recommendations by itself. However, memory metrics can be used by Compute Optimizer to enhance its recommendations if enabled3.
Option E is incorrect because creating an EC2 Instance Savings Plan for the AWS Regions, instance families, and operating systems of interest will not help with cost-optimization and sizing of EC2 instances. Savings Plans are a flexible pricing model that offer lower prices on Amazon EC2 usage in exchange for a commitment to a consistent amount of usage for a 1- or 3-year term4. Savings Plans do not affect the configuration or utilization of EC2 instances.

NEW QUESTION 99
- (Exam Topic 2)
A company has built a high performance computing (HPC) cluster in AWS tor a tightly coupled workload that generates a large number of shared files stored in Amazon EFS. The cluster was performing well when the number of Amazon EC2 instances in the cluster was 100. However, when the company increased the cluster size to 1,000 EC2 instances, overall performance was well below expectations.
Which collection of design choices should a solutions architect make to achieve the maximum performance from the HPC cluster? (Select THREE.)

A. Ensure the HPC cluster Is launched within a single Availability Zone.
B. Launch the EC2 instances and attach elastic network interfaces in multiples of four.
C. Select EC2 Instance types with an Elastic Fabric Adapter (EFA) enabled.
D. Ensure the cluster Is launched across multiple Availability Zones.
E. Replace Amazon EFS with multiple Amazon EBS volumes in a RAID array.
F. Replace Amazon EFS with Amazon FSx for Lustre.

Answer: ACF

Explanation:
* A. High performance computing (HPC) workload cluster should be in a single AZ.
* C. Elastic Fabric Adapter (EFA) is a network device that you can attach to your Amazon EC2 instances to accelerate High Performance Computing (HPC)

* F. Amazon FSx for Lustre - Use it for workloads where speed matters, such as machine learning, high performance computing (HPC), video processing, and financial modeling.
Cluster – packs instances close together inside an Availability Zone. This strategy enables workloads to achieve the low-latency network performance necessary for tightly-coupled node-to-node communication that is typical of HPC applications.
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html

NEW QUESTION 101
- (Exam Topic 2)
A company runs a customer service center that accepts calls and automatically sends all customers a managed, interactive, two-way experience survey by text message.
The applications that support the customer service center run on machines that the company hosts in an on-premises data center. The hardware that the company uses is old, and the company is experiencing downtime with the system. The company wants to migrate the system to AWS to improve reliability.
Which solution will meet these requirements with the LEAST ongoing operational overhead?

A. Use Amazon Connect to replace the old call center hardwar
B. Use Amazon Pinpoint to send text message surveys to customers.
C. Use Amazon Connect to replace the old call center hardwar
D. Use Amazon Simple Notification Service (Amazon SNS) to send text message surveys to customers.
E. Migrate the call center software to Amazon EC2 instances that are in an Auto Scaling grou
F. Use the EC2 instances to send text message surveys to customers.
G. Use Amazon Pinpoint to replace the old call center hardware and to send text message surveys to customers.

Answer: A

Explanation:
Amazon Connect is a cloud-based contact center service that allows you to set up a virtual call center for your business. It provides an easy-to-use interface for managing customer interactions through voice and chat. Amazon Connect integrates with other AWS services, such as Amazon S3 and Amazon Kinesis, to help you collect, store, and analyze customer data for insights into customer behavior and trends. On the other hand, Amazon Pinpoint is a marketing automation and analytics service that allows you to engage with your customers across different channels, such as email, SMS, push notifications, and voice. It helps you create personalized campaigns based on user behavior and enables you to track user engagement and retention. While both services allow you to communicate with your customers, they serve different purposes. Amazon Connect is focused on customer support and service, while Amazon Pinpoint is focused on marketing and engagement.

NEW QUESTION 102
- (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 Regio
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 Regio
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 Regio
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 Regio
H. Change the Route 53 record to geolocation routing to connect to the API

Answer: C

Explanation:
The company should configure a cross-Region read replica for the RDS database in the new Region. The company should change the Route 53 record to latency-based routing to connect to the API Gateway API. This solution will meet the requirements because a cross-Region read replica is a feature that enables you to create a MariaDB, MySQL, Oracle, PostgreSQL, or SQL Server read replica in a different Region from the source DB instance. You can use cross-Region read replicas to improve availability and disaster recovery, scale out globally, or migrate an existing database to a new Region1. By creating a cross-Region read replica for the RDS database in the new Region, the company can have a standby copy of its primary database that can serve read-only traffic from users in Europe. A latency-based routing policy is a feature that enables you to route traffic based on the latency between your users and your resources. You can use latency-based routing to route traffic to the resource that provides the best latency2. By changing the Route 53 record to latency-based routing, the company can minimize latency for users who download reports by connecting them to the API Gateway API in the Region that provides the best response time.
The other options are not correct because:

- Using AWS Database Migration Service (AWS DMS) to replicate the primary database in the original Region to the database in the new Region would not be as cost-effective or simple as using a
cross-Region read replica. AWS DMS is a service that enables you to migrate relational databases, data
warehouses, NoSQL databases, and other types of data stores. You can use AWS DMS to perform one-time migrations or continuous data replication with high availability and consolidate databases into
a petabyte-scale data warehouse3. However, AWS DMS requires more configuration and management than creating a cross-Region read replica, which is fully managed by Amazon RDS. AWS DMS also incurs additional charges for replication instances and tasks.
- Creating an Amazon API Gateway Data API service integration with Amazon Redshift would not help with disaster recovery or minimizing latency. The Data API is a feature that enables you to query your Amazon Redshift cluster using HTTP requests, without needing a persistent connection or a SQL client. It is useful for building applications that interact with Amazon Redshift, but not for replicating or recovering data from an RDS database.
- Creating an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster would not help with disaster recovery or minimizing latency. AWS Data Exchange is a service that makes it easy for AWS customers to exchange data in the cloud. You can use AWS Data Exchange to subscribe to a diverse selection of third-party data products or offer your own data products to other AWS customers. A datashare is a feature that enables you to share live and secure access to your Amazon Redshift data across your accounts or with third parties without copying or moving the underlying data. It is useful for sharing

query results and views with other users, but not for replicating or recovering data from an RDS database.
References:

- https://aws.amazon.com/dms/
- https://docs.aws.amazon.com/redshift/latest/mgmt/data-api.html
- https://aws.amazon.com/data-exchange/
- https://docs.aws.amazon.com/redshift/latest/dg/datashare-overview.html

NEW QUESTION 105
- (Exam Topic 2)
A solutions architect is designing a solution to process events. The solution must have the ability to scale in and out based on the number of events that the solution receives. If a processing error occurs, the event must move into a separate queue for review.
Which solution will meet these requirements?

A. Send event details to an Amazon Simple Notification Service (Amazon SNS) topi
B. Configure an AWS Lambda function as a subscriber to the SNS topic to process the event
C. Add an on-failure destination to the functio
D. Set an Amazon Simple Queue Service (Amazon SQS) queue as the target.
E. Publish events to an Amazon Simple Queue Service (Amazon SQS) queu
F. Create an Amazon EC2 Auto Scaling grou
G. Configure the Auto Scaling group to scale in and out based on the ApproximateAgeOfOldestMessage metric of the queu
H. Configure the application to write failed messages to a dead-letter queue.
I. Write events to an Amazon DynamoDB tabl
J. Configure a DynamoDB stream for the tabl
K. Configure the stream to invoke an AWS Lambda functio
L. Configure the Lambda function to process the events.
M. Publish events to an Amazon EventBridge event bu
N. Create and run an application on an Amazon EC2 instance with an Auto Scaling group that isbehind an Application Load Balancer (ALB). Set the ALB as the event bus targe
O. Configure the event bus to retry event
P. Write messages to a dead-letter queue if the application cannot process the messages.

Answer: A

Explanation:
Amazon Simple Notification Service (Amazon SNS) is a fully managed pub/sub messaging service that enables users to send messages to multiple subscribers1. Users can send event details to an Amazon SNS topic and configure an AWS Lambda function as a subscriber to the SNS topic to process the events. Lambda is a serverless compute service that runs code in response to events and automatically manages the underlying compute resources2. Users can add an on-failure destination to the function and set an Amazon Simple Queue
Service (Amazon SQS) queue as the target. Amazon SQS is a fully managed message queuing service that enables users to decouple and scale microservices, distributed systems, and serverless applications3. This way, if a processing error occurs, the event will move into the separate queue for review.
Option B is incorrect because publishing events to an Amazon SQS queue and creating an Amazon EC2 Auto Scaling group will not have the ability to scale in and out based on the number of events that the solution receives. Amazon EC2 is a web service that provides secure, resizable compute capacity in the cloud. Auto Scaling is a feature that helps users maintain application availability and allows them to scale their EC2 capacity up or down automatically according to conditions they define. However, for this use case, using SQS and EC2 will not take advantage of the serverless capabilities of Lambda and SNS.
Option C is incorrect because writing events to an Amazon DynamoDB table and configuring a DynamoDB stream for the table will not have the ability to move events into a separate queue for review if a processing error occurs. Amazon DynamoDB is a fully managed key-value and document database that delivers single-digit millisecond performance at any scale. DynamoDB Streams is a feature that captures data
modification events in DynamoDB tables. Users can configure the stream to invoke a Lambda function, but they cannot configure an on-failure destination for the function.
Option D is incorrect because publishing events to an Amazon EventBridge event bus and setting an Application Load Balancer (ALB) as the event bus target will not have the ability to move events into a separate queue for review if a processing error occurs. Amazon EventBridge is a serverless event bus service that makes it easy to connect applications with data from a variety of sources. An ALB is a load balancer that distributes incoming application traffic across multiple targets, such as EC2 instances, containers, IP addresses, Lambda functions, and virtual appliances. Users can configure EventBridge to retry events, but they cannot configure an on-failure destination for the ALB.

NEW QUESTION 106
- (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 selectio
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-attribut

NEW QUESTION 109
- (Exam Topic 2)
A company processes environment data. The has a set up sensors to provide a continuous stream of data from different areas in a city. The data is available in JSON format.

The company wants to use an AWS solution to send the data to a database that does not require fixed schemas for storage. The data must be send in real time.
Which solution will meet these requirements?

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

Answer: B

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

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

NEW QUESTION 112
- (Exam Topic 1)
A company has deployed an application on AWS Elastic Beanstalk. The application uses Amazon Aurora for the database layer. An Amazon CloudFront distribution serves web requests and includes the Elastic Beanstalk domain name as the origin server. The distribution is configured with an alternate domain name that visitors use when they access the application.
Each week, the company takes the application out of service for routine maintenance. During the time that the application is unavailable, the company wants visitors to receive an informational message instead of a
CloudFront error message.
A solutions architect creates an Amazon S3 bucket as the first step in the process.
Which combination of steps should the solutions architect take next to meet the requirements? (Choose three.)

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

Answer: ACD

Explanation:
The company wants to serve static content from an S3 bucket during the maintenance period. To do this, the following steps are required:

- Upload static informational content to the S3 bucket. This will provide the source of the content that will be served to the visitors.
- Set the S3 bucket as a second origin in the original CloudFront distribution. Configure the distribution and the S3 bucket to use an origin access identity (OAI). This will allow CloudFront to access the S3 bucket securely and prevent public access to the bucket.
- During the weekly maintenance, edit the default cache behavior to use the S3 origin. Revert the change when the maintenance is complete. This will redirect all web requests to the S3 bucket instead of the Elastic Beanstalk domain name.

The other options are not correct because:

- Creating a new CloudFront distribution is not necessary and would require changing the alternate domain name configuration.
- Creating a cache behavior for the S3 origin on a new distribution would not work because the visitors would still access the original distribution using the alternate domain name.
- Configuring Elastic Beanstalk to serve traffic from the S3 bucket is not possible and would not achieve the desired result.

References:

- https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/distribution-web-values-specify.

NEW QUESTION 116
- (Exam Topic 1)
A company is running an event ticketing platform on AWS and wants to optimize the platform's
cost-effectiveness. The platform is deployed on Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 and is backed by an Amazon RDS for MySQL DB instance. The company is developing new application features to run on Amazon EKS with AWS Fargate.
The platform experiences infrequent high peaks in demand. The surges in demand depend on event dates. Which solution will provide the MOST cost-effective setup for the platform?

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

O. Temporarily scale up the DB instance manually during peaks.

Answer: B

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

NEW QUESTION 119
- (Exam Topic 1)
A company has an on-premises monitoring solution using a PostgreSQL database for persistence of events. The database is unable to scale due to heavy ingestion and it frequently runs out of storage.
The company wants to create a hybrid solution and has already set up a VPN connection between its network and AWS. The solution should include the following attributes:
• Managed AWS services to minimize operational complexity
• A buffer that automatically scales to match the throughput of data and requires no on-going administration.
• A visualization toot to create dashboards to observe events in near-real time.
• Support for semi -structured JSON data and dynamic schemas.
Which combination of components will enabled© company to create a monitoring solution that will satisfy these requirements" (Select TWO.)

A. Use Amazon Kinesis Data Firehose to buffer events Create an AWS Lambda function 10 process and transform events
B. Create an Amazon Kinesis data stream to buffer events Create an AWS Lambda function to process and transform evens
C. Configure an Amazon Aurora PostgreSQL DB cluster to receive events Use Amazon Quick Sight to read from the database and create near-real-time visualizations and dashboards
D. Configure Amazon Elasticsearch Service (Amazon ES) to receive events Use the Kibana endpoint deployed with Amazon ES to create near-real-time visualizations and dashboards.
E. Configure an Amazon Neptune 0 DB instance to receive events Use Amazon QuickSight to read from the database and create near-real-time visualizations and dashboards

Answer: AD

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

NEW QUESTION 122
- (Exam Topic 1)
A company is hosting an image-processing service on AWS in a VPC. The VPC extends across two Availability Zones. Each Availability Zone contains one public subnet and one private subnet.
The service runs on Amazon EC2 instances in the private subnets. An Application Load Balancer in the public subnets is in front of the service. The service needs to communicate with the internet and does so through two NAT gateways. The service uses Amazon S3 for image storage. The EC2 instances retrieve approximately 1 ¢’ of data from an S3 bucket each day.
The company has promoted the service as highly secure. A solutions architect must reduce cloud expenditures as much as possible without compromising the service's security posture or increasing the time spent on ongoing operations.
Which solution will meet these requirements?

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

Answer: C

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

NEW QUESTION 126
.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons 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