

SAP-C02 Dumps

AWS Certified Solutions Architect - Professional

<https://www.certleader.com/SAP-C02-dumps.html>



NEW QUESTION 1

- (Exam Topic 1)

A company is running an application on several Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer. The load on the application varies throughout the day, and EC2 instances are scaled in and out on a regular basis. Log files from the EC2 instances are copied to a central Amazon S3 bucket every 15 minutes. The security team discovers that log files are missing from some of the terminated EC2 instances.

Which set of actions will ensure that log files are copied to the central S3 bucket from the terminated EC2 instances?

- A. Create a script to copy log files to Amazon S3, and store the script in a file on the EC2 instance
- B. Create an Auto Scaling lifecycle hook and an Amazon EventBridge (Amazon CloudWatch Events) rule to detect lifecycle events from the Auto Scaling group
- C. Invoke an AWS Lambda function on the autoscaling:EC2_INSTANCE_TERMINATING transition to send ABANDON to the Auto Scaling group to prevent termination, run the script to copy the log files, and terminate the instance using the AWS SDK.
- D. Create an AWS Systems Manager document with a script to copy log files to Amazon S3. Create an Auto Scaling lifecycle hook and an Amazon EventBridge (Amazon CloudWatch Events) rule to detect lifecycle events from the Auto Scaling group
- E. Invoke an AWS Lambda function on the autoscaling:EC2_INSTANCE_TERMINATING transition to call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send CONTINUE to the Auto Scaling group to terminate the instance.
- F. Change the log delivery rate to every 5 minutes
- G. Create a script to copy log files to Amazon S3, and add the script to EC2 instance user data
- H. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to detect EC2 instance termination
- I. Invoke an AWS Lambda function from the EventBridge (CloudWatch Events) rule that uses the AWS CLI to run the user-data script to copy the log files and terminate the instance.
- J. Create an AWS Systems Manager document with a script to copy log files to Amazon S3. Create an Auto Scaling lifecycle hook that publishes a message to an Amazon Simple Notification Service (Amazon SNS) topic
- K. From the SNS notification, call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send ABANDON to the Auto Scaling group to terminate the instance.

Answer: B

Explanation:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/adding-lifecycle-hooks.html>

- Refer to Default Result section - If the instance is terminating, both abandon and continue allow the instance to terminate. However, abandon stops any remaining actions, such as other lifecycle hooks, and continue allows any other lifecycle hooks to complete.

[https://aws.amazon.com/blogs/infrastructure-and-automation/run-code-before-terminating-an-ec2-auto-scaling-i](https://aws.amazon.com/blogs/infrastructure-and-automation/run-code-before-terminating-an-ec2-auto-scaling-instance/) <https://github.com/aws-samples/aws-lambda-lifecycle-hooks-function>

<https://github.com/aws-samples/aws-lambda-lifecycle-hooks-function/blob/master/cloudformation/template.yaml>

NEW QUESTION 2

- (Exam Topic 1)

An online e-commerce business is running a workload on AWS. The application architecture includes a web tier, an application tier for business logic, and a database tier for user and transactional data management. The database server has a 100 GB memory requirement. The business requires cost-efficient disaster recovery for the application with an RTO of 5 minutes and an RPO of 1 hour. The business also has a regulatory requirement for out-of-region disaster recovery with a minimum distance between the primary and alternate sites of 250 miles.

Which of the following options can the solutions architect design to create a comprehensive solution for this customer that meets the disaster recovery requirements?

- A. Back up the application and database data frequently and copy them to Amazon S3. Replicate the backups using S3 cross-region replication, and use AWS CloudFormation to instantiate infrastructure for disaster recovery and restore data from Amazon S3.
- B. Employ a pilot light environment in which the primary database is configured with mirroring to build a standby database on m4.large in the alternate region
- C. Use AWS CloudFormation to instantiate the web servers, application servers, and load balancers in case of a disaster to bring the application up in the alternate region
- D. Vertically resize the database to meet the full production demands, and use Amazon Route 53 to switch traffic to the alternate region.
- E. Use a scaled-down version of the fully functional production environment in the alternate region that includes one instance of the web server, one instance of the application server, and a replicated instance of the database server in standby mode
- F. Place the web and the application tiers in an Auto Scaling group behind a load balancer, which can automatically scale when the load arrives to the application
- G. Use Amazon Route 53 to switch traffic to the alternate region,
- H. Employ a multi-region solution with fully functional web
- I. application, and database tiers in both regions with equivalent capacity
- J. Activate the primary database in one region only and the standby database in the other region
- K. Use Amazon Route 53 to automatically switch traffic from one region to another using health check routing policies.

Answer: C

Explanation:

As RTO is in minutes

(<https://docs.aws.amazon.com/wellarchitected/latest/reliability-pillar/plan-for-disaster-recovery-dr.html>) Warm standby (RPO in seconds, RTO in minutes): Maintain a scaled-down version of a fully functional environment always running in the DR Region. Business-critical systems are fully duplicated and are always on, but with a scaled-down fleet. When the time comes for recovery, the system is scaled up quickly to handle the production load.

NEW QUESTION 3

- (Exam Topic 1)

A solutions architect is designing the data storage and retrieval architecture for a new application that a company will be launching soon. The application is designed to ingest millions of small records per minute from devices all around the world. Each record is less than 4 KB in size and needs to be stored in a durable location where it can be retrieved with low latency. The data is ephemeral and the company is required to store the data for 120 days only, after which the data can be deleted.

The solutions architect calculates that, during the course of a year, the storage requirements would be about 10-15 TB.

Which storage strategy is the MOST cost-effective and meets the design requirements?

- A. Design the application to store each incoming record as a single .csv file in an Amazon S3 bucket to allow for indexed retrieval
- B. Configure a lifecycle policy to delete data older than 120 days.
- C. Design the application to store each incoming record in an Amazon DynamoDB table properly configured for the scale

- D. Configure the DynamoDB Time to Live (TTL) feature to delete records older than 120 days.
- E. Design the application to store each incoming record in a single table in an Amazon RDS MySQL database.
- F. Run a nightly cron job that executes a query to delete any records older than 120 days.
- G. Design the application to batch incoming records before writing them to an Amazon S3 bucket.
- H. Update the metadata for the object to contain the list of records in the batch and use the Amazon S3 metadata search feature to retrieve the data.
- I. Configure a lifecycle policy to delete the data after 120 days.

Answer: B

Explanation:

DynamoDB with TTL, cheaper for sustained throughput of small items + suited for fast retrievals. S3 cheaper for storage only, much higher costs with writes. RDS not designed for this use case.

NEW QUESTION 4

- (Exam Topic 1)

A solutions architect has an operational workload deployed on Amazon EC2 instances in an Auto Scaling group. The VPC architecture spans two Availability Zones (AZ) with a subnet in each that the Auto Scaling group is targeting. The VPC is connected to an on-premises environment and connectivity cannot be interrupted. The maximum size of the Auto Scaling group is 20 instances in service. The VPC IPv4 addressing is as follows:

VPC CIDR: 10.0.0.0/23

AZ1 subnet CIDR: 10.0.0.0/24 AZ2 subnet CIDR: 10.0.1.0/24

Since deployment, a third AZ has become available in the Region. The solutions architect wants to adopt the new AZ without adding additional IPv4 address space and without service downtime.

Which solution will meet these requirements?

- A. Update the Auto Scaling group to use the AZ2 subnet only.
- B. Delete and re-create the AZ1 subnet using half the previous address space.
- C. Adjust the Auto Scaling group to also use the new AZ1 subnet.
- D. When the instances are healthy, adjust the Auto Scaling group to use the AZ1 subnet only.
- E. Remove the current AZ2 subnet.
- F. Create a new AZ2 subnet using the second half of the address space from the original AZ1 subnet.
- G. Create a new AZ3 subnet using half the original AZ2 subnet address space, then update the Auto Scaling group to target all three new subnets.
- H. Terminate the EC2 instances in the AZ1 subnet.
- I. Delete and re-create the AZ1 subnet using half the address space.
- J. Update the Auto Scaling group to use this new subnet.
- K. Repeat this for the second AZ.
- L. Define a new subnet in AZ3, then update the Auto Scaling group to target all three new subnets.
- M. Create a new VPC with the same IPv4 address space and define three subnets, with one for each AZ.
- N. Update the existing Auto Scaling group to target the new subnets in the new VPC.
- O. Update the Auto Scaling group to use the AZ2 subnet only.
- P. Update the AZ1 subnet to have half the previous address space.
- Q. Adjust the Auto Scaling group to also use the AZ1 subnet again.
- R. When the instances are healthy, adjust the Auto Scaling group to use the AZ1 subnet only.
- S. Update the current AZ2 subnet and assign the second half of the address space from the original AZ1 subnet.
- T. Create a new AZ3 subnet using half the original AZ2 subnet address space, then update the Auto Scaling group to target all three new subnets.

Answer: A

Explanation:

https://aws.amazon.com/premiumsupport/knowledge-center/vpc-ip-address-range/?nc1=h_ls

It's not possible to modify the IP address range of an existing virtual private cloud (VPC) or subnet. You must delete the VPC or subnet, and then create a new VPC or subnet with your preferred CIDR block.

NEW QUESTION 5

- (Exam Topic 1)

A company wants to change its internal cloud billing strategy for each of its business units. Currently, the cloud governance team shares reports for overall cloud spending with the head of each business unit. The company uses AWS Organizations to manage the separate AWS accounts for each business unit. The existing tagging standard in Organizations includes the application, environment, and owner. The cloud governance team wants a centralized solution so each business unit receives monthly reports on its cloud spending. The solution should also send notifications for any cloud spending that exceeds a set threshold.

Which solution is the MOST cost-effective way to meet these requirements?

- A. Configure AWS Budgets in each account and configure budget alerts that are grouped by application, environment, and owner.
- B. Add each business unit to an Amazon SNS topic for each alert.
- C. Use Cost Explorer in each account to create monthly reports for each business unit.
- D. Configure AWS Budgets in the organization's master account and configure budget alerts that are grouped by application, environment, and owner.
- E. Add each business unit to an Amazon SNS topic for each alert.
- F. Use Cost Explorer in the organization's master account to create monthly reports for each business unit.
- G. Configure AWS Budgets in each account and configure budget alerts that are grouped by application, environment, and owner.
- H. Add each business unit to an Amazon SNS topic for each alert.
- I. Use the AWS Billing and Cost Management dashboard in each account to create monthly reports for each business unit.
- J. Enable AWS Cost and Usage Reports in the organization's master account and configure reports grouped by application, environment, and owner.
- K. Create an AWS Lambda function that processes AWS Cost and Usage Reports, sends budget alerts, and sends monthly reports to each business unit's email list.

Answer: B

Explanation:

Configure AWS Budgets in the organization's master account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in the organization's master account to create monthly reports for each business unit.

<https://aws.amazon.com/about-aws/whats-new/2019/07/introducing-aws-budgets-reports/#:~:text=AWS%20Bud>

NEW QUESTION 6

- (Exam Topic 1)

A scientific organization requires the processing of text and picture data stored in an Amazon S3 bucket. The data is gathered from numerous radar stations during a mission's live, time-critical phase. The data is uploaded by the radar stations to the source S3 bucket. The data is preceded with the identification number of the radar station.

In a second account, the business built a destination S3 bucket. To satisfy a compliance target, data must be transferred from the source S3 bucket to the destination S3 bucket. Replication is accomplished by using an S3 replication rule that covers all items in the source S3 bucket.

A single radar station has been recognized as having the most precise data. At this radar station, data replication must be completed within 30 minutes of the radar station uploading the items to the source S3 bucket.

What actions should a solutions architect take to ensure that these criteria are met?

- A. Set up an AWS DataSync agent to replicate the prefixed data from the source S3 bucket to the destination S3 bucket
- B. Select to use at available bandwidth on the task, and monitor the task to ensure that it is in the TRANSFERRING status
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert if this status changes.
- D. In the second account, create another S3 bucket to receive data from the radar station with the most accurate data. Set up a new replication rule for this new S3 bucket to separate the replication from the other radar stations. Monitor the maximum replication time to the destination.
- E. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert when the time exceeds the desired threshold.
- F. Enable Amazon S3 Transfer Acceleration on the source S3 bucket, and configure the radar station with the most accurate data to use the new endpoint. Monitor the S3 destination bucket's TotalRequestLatency metric. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert if this status changes.
- G. Create a new S3 replication rule on the source S3 bucket that filters for the keys that use the prefix of the radar station with the most accurate data. Enable S3 Replication Time Control (S3 RTC). Monitor the maximum replication time to the destination. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert when the time exceeds the desired threshold.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-time-control.html>

NEW QUESTION 7

- (Exam Topic 1)

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

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

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

Answer: AD

Explanation:

https://docs.aws.amazon.com/organizations/latest/APIReference/API_EnableAllFeatures.html

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

NEW QUESTION 8

- (Exam Topic 1)

A large payroll company recently merged with a small staffing company. The unified company now has multiple business units, each with its own existing AWS account.

A solutions architect must ensure that the company can centrally manage the billing and access policies for all the AWS accounts. The solutions architect configures AWS Organizations by sending an invitation to all member accounts of the company from a centralized management account.

What should the solutions architect do next to meet these requirements?

- A. Create the OrganizationAccountAccess IAM group in each member account.
- B. Include the necessary IAM roles for each administrator.
- C. Create the OrganizationAccountAccessPolicy IAM policy in each member account.
- D. Connect the member accounts to the management account by using cross-account access.
- E. Create the OrganizationAccountAccessRole IAM role in each member account.
- F. Grant permission to the management account to assume the IAM role.
- G. Create the OrganizationAccountAccessRole IAM role in the management account. Attach the Administrator Access AWS managed policy to the IAM role.
- H. Assign the IAM role to the administrators in each member account.

Answer: C

NEW QUESTION 9

- (Exam Topic 1)

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

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

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

Answer: B

Explanation:

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

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

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

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

NEW QUESTION 10

- (Exam Topic 1)

A company is running an Apache Hadoop cluster on Amazon EC2 instances. The Hadoop cluster stores approximately 100 TB of data for weekly operational reports and allows occasional access for data scientists to retrieve data. The company needs to reduce the cost and operational complexity for storing and serving this data.

Which solution meets these requirements in the MOST cost-effective manner?

- A. Move the Hadoop cluster from EC2 instances to Amazon EM
- B. Allow data access patterns to remain the same.
- C. Write a script that resizes the EC2 instances to a smaller instance type during downtime and resizes the instances to a larger instance type before the reports are created.
- D. Move the data to Amazon S3 and use Amazon Athena to query the data for report
- E. Allow the data scientists to access the data directly in Amazon S3.
- F. Migrate the data to Amazon DynamoDB and modify the reports to fetch data from DynamoD
- G. Allow the data scientists to access the data directly in DynamoDB.

Answer: C

Explanation:

"The company needs to reduce the cost and operational complexity for storing and serving this data. Which solution meets these requirements in the MOST cost-effective manner?" EMR storage is ephemeral. The company has 100TB that need to persist, they would have to use EMRFS to backup to S3 anyway.

<https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-plan-storage.html>

100TB

EBS - 8.109\$ S3 - 2.355\$

You have saved 5.752\$

This amount can be used for Athen. BTW. we don't know indexes, amount of data that is scanned. What we know is that it will be: "occasional access for data scientists to retrieve data"

NEW QUESTION 10

- (Exam Topic 1)

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

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

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

Answer: AE

NEW QUESTION 12

- (Exam Topic 1)

A media company uses Amazon DynamoDB to store metadata for its catalog of movies that are available to stream. Each media item Contains user-facing content that concludes a description of the media, a list of search tags, and similar data. In addition, media items include a list of Amazon S3 key names that relate to movie files. The company stores these movie files in a single S3 bucket that has versioning enable. The company uses Amazon CloudFront to serve these movie files.

The company has 100.000 media items, and each media item can have many different S3 objects that represent different encodings of the same media S3 objects that belong to the same media item are grouped together under the same key prefix, which is a random unique ID

Because of an expiring contract with a media provider, the company must remove 2.000 media items. The company must completely delete all DynamoDB keys and movie files on Amazon S3 that are related to these media items within 36 hours The company must ensure that the content cannot be recovered.

Which combination of actions will meet these requirements? (Select TWO.)

- A. Configure the dynamoDB table with a TTL field
- B. Create and invoke an AWS Lambda function to perform a conditional update Set the TTL field to the time of the contract's expiration on every affected media item.
- C. Configure an S3 Lifecycle object expiration rule that is based on the contract's expiration date
- D. Write a script to perform a conditional delete on all the affected DynamoDB records
- E. Temporarily suspend versioning on the S3 bucket
- F. Create and invoke an AWS Lambda function that deletes affected objects Reactivate versioning when the operation is complete
- G. Write a script to delete objects from Amazon S3 Specify in each request a NoncurrentVersionExpiration property with a NoncurrentDays attribute set to 0.

Answer: CE

NEW QUESTION 15

- (Exam Topic 1)

To abide by industry regulations, a solutions architect must design a solution that will store a company's critical data in multiple public AWS Regions, including in the United States, where the company's headquarters is located. The solutions architect is required to provide access to the data stored in AWS to the company's global WAN network. The security team mandates that no traffic accessing this data should traverse the public internet.

How should the solutions architect design a highly available solution that meets the requirements and is cost-effective?

- A. Establish AWS Direct Connect connections from the company headquarters to all AWS Regions in use. Use the company WAN to send traffic over to the headquarters and then to the respective DX connection to access the data.
- B. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connection.
- C. Use inter-region VPC peering to access the data in other AWS Regions.
- D. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connection.
- E. Use an AWS transit VPC solution to access data in other AWS Regions.
- F. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connection.
- G. Use Direct Connect Gateway to access data in other AWS Regions.

Answer: D

Explanation:

This feature also allows you to connect to any of the participating VPCs from any Direct Connect location, further reducing your costs for making using AWS services on a cross-region basis. <https://aws.amazon.com/blogs/aws/new-aws-direct-connect-gateway-inter-region-vpc-access/>
<https://docs.aws.amazon.com/whitepapers/latest/aws-vpc-connectivity-options/aws-direct-connect-aws-transit-g>

NEW QUESTION 16

- (Exam Topic 1)

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

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

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

Answer: B

Explanation:

"This policy uses the Deny effect to deny access to all requests for operations that don't target one of the two approved regions (eu-central-1 and eu-west-1)."
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_general.htm
https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements_condition.html

NEW QUESTION 17

- (Exam Topic 1)

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

Which solution meets the requirements?

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

Answer: C

Explanation:

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

NEW QUESTION 18

- (Exam Topic 1)

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

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

What should the solutions architect do to meet these requirements?

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

Answer: D

Explanation:

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

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

NEW QUESTION 21

- (Exam Topic 1)

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

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

Answer: D

NEW QUESTION 22

- (Exam Topic 1)

A solutions architect is building a web application that uses an Amazon RDS for PostgreSQL DB instance The DB instance is expected to receive many more reads than writes. The solutions architect needs to ensure that the large amount of read traffic can be accommodated and that the DB instance is highly available. Which steps should the solutions architect take to meet these requirements? (Select THREE)

- A. Create multiple read replicas and put them into an Auto Scaling group.
- B. Create multiple read replicas in different Availability Zones.
- C. Create an Amazon Route 53 hosted zone and a record set for each read replica with a TTL and a weighted routing policy.
- D. Create an Application Load Balancer (ALB) and put the read replicas behind the ALB.
- E. Configure an Amazon CloudWatch alarm to detect a failed read replic
- F. Set the alarm to directly invoke an AWS Lambda function to delete its Route 53 record set.
- G. Configure an Amazon Route 53 health check for each read replica using its endpoint

Answer: BCF

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/requests-rds-read-replicas/>

You can use Amazon Route 53 weighted record sets to distribute requests across your read replicas. Within a Route 53 hosted zone, create individual record sets for each DNS endpoint associated with your read replicas and give them the same weight. Then, direct requests to the endpoint of the record set. You can incorporate Route 53 health checks to be sure that Route 53 directs traffic away from unavailable read replicas

NEW QUESTION 24

- (Exam Topic 1)

A medical company is running a REST API on a set of Amazon EC2 instances. The EC2 instances run in an Auto Scaling group behind an Application Load Balancer (ALB). The ALB runs in three public subnets, and the EC2 instances run in three private subnets. The company has deployed an Amazon CloudFront distribution that has the ALB as the only origin.

Which solution should a solutions architect recommend to enhance the origin security?

- A. Store a random string in AWS Secrets Manage
- B. Create an AWS Lambda (unction for automatic secret rotatio
- C. Configure CloudFront to inject the random string as a custom HTTP header for the origin reques
- D. Create an AWS WAF web ACL rule with a string match rule for the custom heade
- E. Associate the web ACL with the ALB.
- F. Create an AWS WAF web ACL rule with an IP match condition of the CloudFront service IP address range
- G. Associate the web ACL with the AL
- H. Move the ALB into the three private subnets.
- I. Store a random string in AWS Systems Manager Parameter Stor
- J. Configure Parameter Store automatic rotation for the strin
- K. Configure CloudFront to inject the random siring as a custom HTTP header for the origin reques
- L. Inspect the value of the custom HTTP header, and block access in the ALB.
- M. Configure AWS Shield Advance

- N. Create a security group policy to allow connections from CloudFront service IP address range
O. Add the policy to AWS Shield Advanced, and attach the policy to the ALB.

Answer: D

Explanation:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-suspend-resume-processes.html>

it shows For Amazon EC2 Auto Scaling, there are two primary process types: Launch and Terminate. The Launch process adds a new Amazon EC2 instance to an Auto Scaling group, increasing its capacity. The Terminate process removes an Amazon EC2 instance from the group, decreasing its capacity. HealthCheck process for EC2 autoscaling is not a primary process! It is a process along with the following AddToLoadBalancer AlarmNotification AZRebalance HealthCheck InstanceRefresh ReplaceUnhealthy ScheduledActions From the requirements, Some EC2 instances are now being marked as unhealthy and are being terminated. Application is running at reduced capacity not because instances are marked unhealthy but because they are being terminated.

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-suspend-resume-processes.html#choosing-suspend-r>

NEW QUESTION 29

- (Exam Topic 1)

A company is moving a business-critical multi-tier application to AWS. The architecture consists of a desktop client application and server infrastructure. The server infrastructure resides in an on-premises data center that frequently fails to maintain the application uptime SLA of 99.95%. A solutions architect must re-architect the application to ensure that it can meet or exceed the SLA.

The application contains a PostgreSQL database running on a single virtual machine. The business logic and presentation layers are load balanced between multiple virtual machines. Remote users complain about slow load times while using this latency-sensitive application.

Which of the following will meet the availability requirements with little change to the application while improving user experience and minimizing costs?

- A. Migrate the database to a PostgreSQL database in Amazon EC2. Host the application and presentation layers in automatically scaled Amazon ECS containers behind an Application Load Balance
B. Allocate an Amazon Workspaces Workspace for each end user to improve the user experience.
C. Migrate the database to an Amazon RDS Aurora PostgreSQL configuratio
D. Host the application and presentation layers in an Auto Scaling configuration on Amazon EC2 instances behind an Application Load Balance
E. Use Amazon AppStream 2.0 to improve the user experience.
F. Migrate the database to an Amazon RDS PostgreSQL Multi-AZ configuratio
G. Host the application and presentation layers in automatically scaled AWS Fargate containers behind a Network Load Balance
H. Use Amazon ElastiCache to improve the user experience.
I. Migrate the database to an Amazon Redshift cluster with at least two node
J. Combine and host the application and presentation layers in automatically scaled Amazon ECS containers behind an Application Load Balance
K. Use Amazon CloudFront to improve the user experience.

Answer: B

Explanation:

Aurora would improve availability that can replicate to multiple AZ (6 copies). Auto scaling would improve the performance together with a ALB. AppStream is like Citrix that deliver hosted Apps to users.

NEW QUESTION 30

- (Exam Topic 1)

A company needs to run a software package that has a license that must be run on the same physical host for the duration of its use. The software package is only going to be used for 90 days The company requires patching and restarting of all instances every 30 days

How can these requirements be met using AWS?

- A. Run a dedicated instance with auto-placement disabled.
B. Run the instance on a dedicated host with Host Affinity set to Host.
C. Run an On-Demand Instance with a Reserved Instance to ensure consistent placement.
D. Run the instance on a licensed host with termination set for 90 days.

Answer: B

Explanation:

Host Affinity is configured at the instance level. It establishes a launch relationship between an instance and a Dedicated Host. (This set which host the instance can run on) Auto-placement allows you to manage whether instances that you launch are launched onto a specific host, or onto any available host that has matching configurations. Auto-placement must be configured at the host level. (This sets which instance the host can run.) When affinity is set to Host, an instance launched onto a specific host always restarts on the same host if stopped. This applies to both targeted and untargeted launches.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/how-dedicated-hosts-work.html>

When affinity is set to Off, and you stop and restart the instance, it can be restarted on any available host. However, it tries to launch back onto the last Dedicated Host on which it ran (on a best-effort basis).

NEW QUESTION 32

- (Exam Topic 1)

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 36

- (Exam Topic 1)

A company is using AWS CodePipeline for the CI/CO of an application to an Amazon EC2 Auto Scaling group. All AWS resources are defined in AWS CloudFormation templates. The application artifacts are stored in an Amazon S3 bucket and deployed to the Auto Scaling group using instance user data scripts. As the application has become more complex, recent resource changes in the Cloud Formation templates have caused unplanned downtime. How should a solutions architect improve the CI/CD pipeline to reduce the likelihood that changes in the templates will cause downtime?

- A. Adapt the deployment scripts to detect and report CloudFormation error conditions when performing deployment
- B. Write test plans for a testing team to execute in a non-production environment before approving the change for production.
- C. Implement automated testing using AWS CodeBuild in a test environment
- D. Use CloudFormation changesets to evaluate changes before deployment
- E. Use AWS CodeDeploy to leverage blue/green deployment patterns to allow evaluations and the ability to revert changes, if needed.
- F. Use plugins for the integrated development environment (IDE) to check the templates for errors, and use the AWS CLI to validate that the templates are correct
- G. Adapt the deployment code to check for error conditions and generate notifications on error
- H. Deploy to a test environment and execute a manual test plan before approving the change for production.
- I. Use AWS CodeDeploy and a blue/green deployment pattern with CloudFormation to replace the user data deployment script
- J. Have the operators log in to running instances and go through a manual test plan to verify the application is running as expected.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/devops/performing-bluegreen-deployments-with-aws-codedeploy-and-auto-scaling/> When one adopts go infrastructure as code, we need to test the infrastructure code as well via automated testing, and revert to original if things are not performing correctly.

NEW QUESTION 38

- (Exam Topic 1)

A multimedia company needs to deliver its video-on-demand (VOD) content to its subscribers in a cost-effective way. The video files range in size from 1-15 GB and are typically viewed frequently for the first 6 months after creation, and then access decreases considerably. The company requires all video files to remain immediately available for subscribers. There are now roughly 30,000 files, and the company anticipates doubling that number over time.

What is the MOST cost-effective solution for delivering the company's VOD content?

- A. Store the video files in an Amazon S3 bucket using S3 Intelligent-Tiering
- B. Use Amazon CloudFront to deliver the content with the S3 bucket as the origin.
- C. Use AWS Elemental MediaConvert and store the adaptive bitrate video files in Amazon S3. Configure an AWS Elemental MediaPackage endpoint to deliver the content from Amazon S3.
- D. Store the video files in Amazon Elastic File System (Amazon EFS) Standard
- E. Enable EFS lifecycle management to move the video files to EFS Infrequent Access after 6 months
- F. Create an Amazon EC2 Auto Scaling group behind an Elastic Load Balancer to deliver the content from Amazon EFS.
- G. Store the video files in Amazon S3 Standard
- H. Create S3 Lifecycle rules to move the video files to S3 Standard-Infrequent Access (S3 Standard-IA) after 6 months and to S3 Glacier Deep Archive after 1 year
- I. Use Amazon CloudFront to deliver the content with the S3 bucket as the origin.

Answer: A

Explanation:

<https://d1.awsstatic.com/whitepapers/amazon-cloudfront-for-media.pdf> <https://aws.amazon.com/solutions/implementations/video-on-demand-on-aws/>

NEW QUESTION 39

- (Exam Topic 1)

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

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

What should the solutions architect do to meet these requirements?

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

Answer: D

Explanation:

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

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

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

NEW QUESTION 40

- (Exam Topic 1)

A company has an internal application running on AWS that is used to track and process shipments in the company's warehouse. Currently, after the system

receives an order, it emails the staff the information needed to ship a package. Once the package is shipped, the staff replies to the email and the order is marked as shipped.

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

- A. Use AWS Batch to configure the different tasks required to ship a package
- B. Have AWS Batch trigger an AWS Lambda function that creates and prints a shipping label
- C. Once that label is scanned
- D. as it leaves the warehouse, have another Lambda function move the process to the next step in the AWS Batch job.
- E. When a new order is created, store the order information in Amazon SQS
- F. Have AWS Lambda check the queue every 5 minutes and process any needed work
- G. When an order needs to be shipped, have Lambda print the label in the warehouse
- H. Once the label has been scanned, as it leaves the warehouse, have an Amazon EC2 instance update Amazon S3.
- I. Update the application to store new order information in Amazon DynamoDB
- J. When a new order is created, trigger an AWS Step Functions workflow, mark the orders as "in progress," and print a package label to the warehouse
- K. Once the label has been scanned and fulfilled, the application will trigger an AWS Lambda function that will mark the order as shipped and complete the workflow.
- L. Store new order information in Amazon EFS
- M. Have instances pull the new information from the NFS and send that information to printers in the warehouse
- N. Once the label has been scanned, as it leaves the warehouse, have Amazon API Gateway call the instances to remove the order information from Amazon EFS.

Answer: C

NEW QUESTION 45

- (Exam Topic 1)

A company has a policy that all Amazon EC2 instances that are running a database must exist within the same subnets in a shared VPC. Administrators must follow security compliance requirements and are not allowed to directly log in to the shared account. All company accounts are members of the same organization in AWS Organizations. The number of accounts will rapidly increase as the company grows.

A solutions architect uses AWS Resource Access Manager to create a resource share in the shared account. What is the MOST operationally efficient configuration to meet these requirements?

- A. Add the VPC to the resource share
- B. Add the account IDs as principals
- C. Add all subnets within the VPC to the resource share
- D. Add the account IDs as principals
- E. Add all subnets within the VPC to the resource share
- F. Add the organization as a principal.
- G. Add the VPC to the resource share
- H. Add the organization as a principal

Answer: C

Explanation:

<https://docs.aws.amazon.com/ram/latest/userguide/getting-started-sharing.html#getting-started-sharing-create> To restrict resource sharing to only principals in your organization, choose Allow sharing with principals in your organization only.

<https://docs.aws.amazon.com/ram/latest/userguide/ram-ug.pdf>

NEW QUESTION 50

- (Exam Topic 1)

A company is migrating its three-tier web application from on-premises to the AWS Cloud. The company has the following requirements for the migration process:

- Ingest machine images from the on-premises environment.
- Synchronize changes from the on-premises environment to the AWS environment until the production cutover.
- Minimize downtime when executing the production cutover.
- Migrate the virtual machines' root volumes and data volumes.

Which solution will satisfy these requirements with minimal operational overhead?

- A. Use AWS Server Migration Service (SMS) to create and launch a replication job for each tier of the application
- B. Launch instances from the AMIs created by AWS SMS
- C. After initial testing, perform a final replication and create new instances from the updated AMIs.
- D. Create an AWS CLI VM Import/Export script to migrate each virtual machine
- E. Schedule the script to run incrementally to maintain changes in the application
- F. Launch instances from the AMIs created by VM Import/Export
- G. Once testing is done, rerun the script to do a final import and launch the instances from the AMIs.
- H. Use AWS Server Migration Service (SMS) to upload the operating system volume
- I. Use the AWS CLI import-snapshots command for the data volume
- J. Launch instances from the AMIs created by AWS SMS and attach the data volumes to the instance
- K. After initial testing, perform a final replication, launch new instances from the replicated AMI
- L. and attach the data volumes to the instances.
- M. Use AWS Application Discovery Service and AWS Migration Hub to group the virtual machines as an application
- N. Use the AWS CLI VM Import/Export script to import the virtual machines as AMI
- O. Schedule the script to run incrementally to maintain changes in the application
- P. Launch instances from the AMI
- Q. After initial testing, perform a final virtual machine import and launch new instances from the AMIs.

Answer: A

Explanation:

SMS can handle migrating the data volumes:

<https://aws.amazon.com/about-aws/whats-new/2018/09/aws-server-migration-service-adds-support-for-migrating>

NEW QUESTION 54

- (Exam Topic 1)

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

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

Which migration strategy meets these requirements?

- A. Create an AWS Database Migration Service (AWS DMS) replication instance.
- B. Authorize the public IP address of the replication instance to reach the data warehouse through the corporate firewall. Create a migration task to run at the beginning of the data freeze period.
- C. Install the AWS SCT extraction agents on the on-premises server.
- D. Define the extract, upload, and copy tasks to send the data to an Amazon S3 bucket.
- E. Copy the data into the Amazon Redshift cluster.
- F. Run the tasks at the beginning of the data freeze period.
- G. Install the AWS SCT extraction agents on the on-premises server.
- H. Create a Site-to-Site VPN connection. Create an AWS Database Migration Service (AWS DMS) replication instance that is the appropriate size. Authorize the IP address of the replication instance to be able to access the on-premises data warehouse through the VPN connection.
- I. Create a job in AWS Snowball Edge to import data into Amazon S3. Install AWS SCT extraction agents on the on-premises servers. Define the local and AWS Database Migration Service (AWS DMS) tasks to send the data to the Snowball Edge device. When the Snowball Edge device is returned to AWS and the data is available in Amazon S3, run the AWS DMS subtask to copy the data to Amazon Redshift.

Answer: D

Explanation:

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

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

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

NEW QUESTION 57

- (Exam Topic 1)

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

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

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 58

- (Exam Topic 1)

A company is running an application on Amazon EC2 instances in three environments; development, testing, and production. The company uses AMIs to deploy the EC2 instances. The company builds the AMIs by using custom deployment scripts and infrastructure orchestration tools for each release in each environment. The company is receiving errors in its deployment process. Errors appear during operating system package downloads and during application code installation from a third-party Git hosting service. The company needs deployments to become more reliable across all environments.

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

- A. Mirror the application code to an AWS CodeCommit Git repository.
- B. Use the repository to build EC2 AMIs.
- C. Produce multiple EC2 AMIs.
- D. One for each environment, for each release.
- E. Produce one EC2 AMI for each release for use across all environments.
- F. Mirror the application code to a third-party Git repository that uses Amazon S3 storage.
- G. Use the repository for deployment.
- H. Replace the custom scripts and tools with AWS CodeBuild.
- I. Update the infrastructure deployment process to use EC2 Image Builder.

Answer: ACE

NEW QUESTION 63

- (Exam Topic 2)

A company wants to allow its marketing team to perform SQL queries on customer records to identify market segments. The data is spread across hundreds of files. The records must be encrypted in transit and at rest. The team manager must have the ability to manage users and groups but no team members should have access to services or resources not required for the SQL queries. Additionally, administrators need to audit the queries made and receive notifications when a query violates rules defined by the security team.

AWS Organizations has been used to create a new account and an AWS IAM user with administrator permissions for the team manager. Which design meets these requirements?

- A. Apply a service control policy (SCP) that allows access to IAM Amazon RD
- B. and AWS CloudTrail Load customer records in Amazon RDS MySQL and train users to run queries using the AWS CL
- C. Stream the query logs to Amazon CloudWatch Logs from the RDS database instance Use a subscription filter with AWS Lambda functions to audit and alarm on queries against personal data
- D. Apply a service control policy (SCP) that denies access to all services except IAM Amazon Athena Amazon S3 and AWS CloudTrail Store customer record files in Amazon S3 and tram users to run queries using the CLI via Athena Analyze CloudTrail events to audit and alarm on queries against personal data
- E. Apply a service control policy (SCP) that denies access to all services except IAM Amazon DynamoD
- F. and AWS CloudTrail Store customer records in DynamoDB and train users to run queries using the AWS CLI Enable DynamoDB streams to track the queries that are issued and use an AWS Lambda function for real-time monitoring and alerting
- G. Apply a service control policy (SCP) that allows access to IAM Amazon Athena; Amazon S3, and AWS CloudTrail Store customer records as files in Amazon S3 and train users to leverage the Amazon S3 Select feature and run queries using the AWS CLI Enable S3 object-level logging and analyze CloudTrail events to audit and alarm on queries against personal data

Answer: B

NEW QUESTION 68

- (Exam Topic 2)

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

Which strategy should the solutions architect use?

- A. Use AWS Lambda to run the applicatio
- B. Use Amazon CloudWatch Logs to invoke the Lambda function every 4 hours
- C. Use AWS Batch to run the application Use an AWS Step Functions state machine to invoke the AWS Batch job every 4 hours
- D. Use AWS Fargate to run the application Use Amazon EventBridge (Amazon CloudWatch Events) to invoke the Fargate task every 4 hours
- E. Use Amazon 6C2 Spot Instances to run the application Use AWS CodeDeploy to deploy and run the application every 4 hours.

Answer: C

NEW QUESTION 70

- (Exam Topic 2)

A company is running a three-tier web application in an on-premises data center. The frontend is served by an Apache web server, the middle tier is a monolithic Java application, and the storage tier is a PostgreSQL database.

During a recent marketing promotion, customers could not place orders through the application because the application crashed An analysis showed that all three tiers were overloaded. The application became unresponsive, and the database reached its capacity limit because of read operations. The company already has several similar promotions scheduled in the near future.

A solutions architect must develop a plan for migration to AWS to resolve these issues. The solution must maximize scalability and must minimize operational effort.

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

- A. Refactor the frontend so that static assets can be hosted on Amazon S3. Use Amazon CloudFront to serve the frontend to customer
- B. Connect the frontend to the Java application.
- C. Rehost the Apache web server of the frontend on Amazon EC2 instances that are in an Auto Scaling grou
- D. Use a load balancer in front of the Auto Scaling grou
- E. Use Amazon Elastic File System (Amazon EFS) to host the static assets that the Apache web server needs.
- F. Rehost the Java application in an AWS Elastic Beanstalk environment that includes auto scaling.
- G. Refactor the Java applicatio
- H. Develop a Docker container to run the Java applicatio
- I. Use AWS Fargate to host the container.
- J. Use AWS Database Migration Service (AWS DMS) to replatform the PostgreSQL database to an Amazon Aurora PostgreSQL databas
- K. Use Aurora Auto Scaling for read replicas.
- L. Rehost the PostgreSQL database on an Amazon EC2 instance that has twice as much memory as the on-premises server.

Answer: BCF

NEW QUESTION 72

- (Exam Topic 2)

A company is migrating an on-premises application and a MySQL database to AWS. The application processes highly sensitive data, and new data is constantly updated in the database. The data must not be transferred over the internet. The company also must encrypt the data in transit and at rest.

The database is 5 TB in size. The company already has created the database schema in an Amazon RDS for MySQL DB instance The company has set up a 1 Gbps AWS Direct Connect connection to AWS. The company also has set up a public VIF and a private VIF. A solutions architect needs to design a solution that will migrate the data to AWS with the least possible downtime

Which solution will meet these requirements?

- A. Perform a database backu
- B. Copy the backup files to an AWS Snowball Edge Storage Optimized device.Import the backup to Amazon S3. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3) for encryption at rest Use TLS for encryption in transit Import the data from Amazon S3 to the DB instance.
- C. Use AWS Database Migration Service (AWS DMS) to migrate the data to AW
- D. Create a DMS replication instance in a private subne
- E. Create VPC endpoints for AWS DM

- F. Configure a DMS task to copy data from the on-premises database to the DB instance by using full load plus change data capture (CDC). Use the AWS Key Management Service (AWS KMS) default key for encryption at rest.
- G. Use TLS for encryption in transit.
- H. Perform a database backup.
- I. Use AWS DataSync to transfer the backup files to Amazon S3. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3) for encryption at rest.
- J. Use TLS for encryption in transit. Import the data from Amazon S3 to the DB instance.
- K. Use Amazon S3 File Gateway. Set up a private connection to Amazon S3 by using AWS PrivateLink. Perform a database backup.
- L. Copy the backup files to Amazon S3. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3) for encryption at rest.
- M. Use TLS for encryption in transit.
- N. Import the data from Amazon S3 to the DB instance.

Answer: D

NEW QUESTION 74

- (Exam Topic 2)

A company has a new security policy. The policy requires the company to log any event that retrieves data from Amazon S3 buckets. The company must save these audit logs in a dedicated S3 bucket. The company created the audit logs S3 bucket in an AWS account that is designated for centralized logging. The S3 bucket has a bucket policy that allows write-only cross-account access. A solutions architect must ensure that all S3 object-level access is being logged for current S3 buckets and future S3 buckets. Which solution will meet these requirements?

- A. Enable server access logging for all current S3 buckets.
- B. Use the audit logs S3 bucket as a destination for audit logs.
- C. Enable replication between all current S3 buckets and the audit logs S3 bucket. Enable S3 Versioning in the audit logs S3 bucket.
- D. Configure S3 Event Notifications for all current S3 buckets to invoke an AWS Lambda function every time objects are accessed. Store Lambda logs in the audit logs S3 bucket.
- E. Enable AWS CloudTrail.
- F. Use the audit logs S3 bucket to store logs. Enable data event logging for S3 event sources, current S3 buckets, and future S3 buckets.

Answer: D

NEW QUESTION 78

- (Exam Topic 2)

A development team is deploying new APIs as serverless applications within a company. The team is currently using the AWS Management Console to provision Amazon API Gateway, AWS Lambda, and Amazon DynamoDB resources. A solutions architect has been tasked with automating the future deployments of these serverless APIs. How can this be accomplished?

- A. Use AWS CloudFormation with a Lambda-backed custom resource to provision API Gateway. Use the `MyS3::DynamoDB::Table` and `AWS::Lambda::Function` resources to create the Amazon DynamoDB table and Lambda functions. Write a script to automate the deployment of the CloudFormation template.
- B. Use the AWS Serverless Application Model to define the resources. Upload a YAML template and application files to the code repository. Use AWS CodePipeline to connect to the code repository and to create an action to build using AWS CodeBuild.
- C. Use the AWS CloudFormation deployment provider in CodePipeline to deploy the solution.
- D. Use AWS CloudFormation to define the serverless application.
- E. Implement versioning on the Lambda functions and create aliases to point to the version.
- F. When deploying, configure weights to implement shifting traffic to the newest version, and gradually update the weights as traffic moves over.
- G. Commit the application code to the AWS CodeCommit code repository.
- H. Use AWS CodePipeline and connect to the CodeCommit code repository. Use AWS CodeBuild to build and deploy the Lambda functions using AWS CodeDeploy. Specify the deployment preference type in CodeDeploy to gradually shift traffic over to the new version.

Answer: B

NEW QUESTION 82

- (Exam Topic 2)

A company has a web application that allows users to upload short videos. The videos are stored on Amazon EBS volumes and analyzed by custom recognition software for categorization.

The website contains static content that has variable traffic with peaks in certain months. The architecture consists of Amazon EC2 instances running in an Auto Scaling group for the web application and EC2 instances running in an Auto Scaling group to process an Amazon SQS queue. The company wants to re-architect the application to reduce operational overhead using AWS managed services where possible and remove dependencies on third-party software. Which solution meets these requirements?

- A. Use Amazon ECS containers for the web application and Spot Instances for the Auto Scaling group that processes the SQS queue.
- B. Replace the custom software with Amazon Rekognition to categorize the videos.
- C. Store the uploaded videos on Amazon EFS and mount the file system to the EC2 instances for the web application.
- D. Process the SQS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos.
- E. Host the web application in Amazon S3. Store the uploaded videos in Amazon S3. Use S3 event notifications to publish events to the SQS queue. Process the SQS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos.
- F. Use AWS Elastic Beanstalk to launch EC2 instances in an Auto Scaling group for the web application and launch a worker environment to process the SQS queue. Replace the custom software with Amazon Rekognition to categorize the videos.

Answer: D

NEW QUESTION 84

- (Exam Topic 2)

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

Which tactics should the Solutions Architect use to achieve these objectives?

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

Answer: B

Explanation:

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

NEW QUESTION 85

- (Exam Topic 2)

A solutions architect uses AWS Organizations to manage several AWS accounts for a company. The full Organizations feature set is activated for the organization. All production AWS accounts exist under an OU that is named "production". Systems operators have full administrative privileges within these accounts by using IAM roles.

The company wants to ensure that security groups in all production accounts do not allow inbound traffic for TCP port 22. All noncompliant security groups must be remediated immediately, and no new rules that allow port 22 can be created.

Which solution will meet these requirements?

- A. Write an SCP that denies the CreateSecurityGroup action with a condition of ec2:ingress rule with value 22. Apply the SCP to the 'production' OU.
- B. Configure an AWS CloudTrail trail for all accounts. Send CloudTrail logs to an Amazon S3 bucket in the Organizations management account.
- C. Configure an AWS Lambda function on the management account with permissions to assume a role in all production accounts to describe and modify security groups.
- D. Configure Amazon S3 to invoke the Lambda function on every PutObject event on the S3 bucket. Configure the Lambda function to analyze each CloudTrail event for noncompliant security group actions and to automatically remediate any issues.
- E. Create an Amazon EventBridge (Amazon CloudWatch Events) event bus in the Organizations management account.
- F. Create an AWS CloudFormation template to deploy configurations that send CreateSecurityGroup events to the event bus from all production accounts. Configure an AWS Lambda function in the management account with permissions to assume a role in all production accounts to describe and modify security groups.
- G. Configure the event bus to invoke the Lambda function. Configure the Lambda function to analyze each event for noncompliant security group actions and to automatically remediate any issues.
- H. Create an AWS CloudFormation template to turn on AWS Config. Activate the INCOMING_SSH_DISABLED AWS Config managed rule. Deploy an AWS Lambda function that will run based on AWS Config findings and will remediate noncompliant resources. Deploy the CloudFormation template by using a StackSet that is assigned to the "production" OU.
- I. Apply an SCP to the OU to deny modification of the resources that the CloudFormation template provisions.

Answer: D

NEW QUESTION 90

- (Exam Topic 2)

A company is running an application in the AWS Cloud. The application runs on containers in an Amazon Elastic Container Service (Amazon ECS) cluster. The ECS tasks use the Fargate launch type. The application's data is relational and is stored in Amazon Aurora MySQL. To meet regulatory requirements, the application must be able to recover to a separate AWS Region in the event of an application failure. In case of a failure, no data can be lost. Which solution will meet these requirements with the LEAST amount of operational overhead?

- A. Provision an Aurora Replica in a different Region.
- B. Set up AWS DataSync for continuous replication of the data to a different Region.
- C. Set up AWS Database Migration Service (AWS DMS) to perform a continuous replication of the data to a different Region.
- D. Use Amazon Data Lifecycle Manager (Amazon DLM) to schedule a snapshot every 5 minutes.

Answer: B

NEW QUESTION 95

- (Exam Topic 2)

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

What should a solution architect do to meet these requirements?

- A. Create an AWS Site-to-Site VPN connection Configure integration between a VPN and AD D
- B. Use an Amazon Workspaces client with MFA support enabled to establish a VPN connection.
- C. Create an AWS Client VPN endpoint Create an AD Connector directory for integration with AD DS Enable MFA for AD Connector Use AWS Client VPN to establish a VPN connection.
- D. Create multiple AWS Site-to-Site VPN connections by using AWS VPN CloudHub Configure integration between AWS VPN CloudHub and AD DS Use AWS Cop4ot to establish a VPN connection.
- E. Create an Amazon WorkLink endpoint Configure integration between Amazon WorkLink and AD D
- F. Enable MFA in Amazon WorkLink Use AWS Client VPN to establish a VPN connection.

Answer: B

NEW QUESTION 100

- (Exam Topic 2)

A solutions architect needs to provide AWS Cost and Usage Report data from a company's AWS Organizations management account The company already has an Amazon S3 bucket to store the reports The reports must be automatically ingested into a database that can be visualized with other tools.

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

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that a new object creation in the S3 bucket will trigger
- B. Create an AWS Cost and Usage Report configuration to deliver the data into the S3 bucket
- C. Configure an AWS Glue crawler that a new object creation in the S3 bucket will trigger.
- D. Create an AWS Lambda function that a new object creation in the S3 bucket will trigger
- E. Create an AWS Glue crawler that me AWS Lambda function will trigger to crawl objects in me S3 bucket
- F. Create an AWS Glue crawler that the Amazon EventBridge (Amazon CloudWatCh Events) rule will trigger to crawl objects m the S3 bucket

Answer: BDF

NEW QUESTION 102

- (Exam Topic 2)

A company is running multiple workloads in the AWS Cloud. The company has separate units for software development The company uses AWS Organizations and federation with SAML to give permissions to developers to manage resources in their AWS accounts The development units each deploy their production workloads into a common production account

Recently, an incident occurred in the production account in which members of a development unit terminated an EC2 instance that belonged to a different development unit. A solutions architect must create a solution that prevents a similar incident from happening in the future. The solution also must a low developers the possibility to manage the instances used for their workloads.

Which strategy will meet these requirements?

- A. Create separate OUs in AWS Organizations for each development unit Assign the created OUs to the company AWS accounts Create separate SCPs with a deny action and a StringNotEquals condition for the DevelopmentUnit resource tag that matches the development unit name Assign the SCP to the corresponding OU
- B. Pass an attribute for DevelopmentUnit as an AWS Security Token Service (AWS STS) session tag during SAML federation Update the IAM policy for the developers' assumed IAM role with a deny action and a StringNotEquals condition for the DevelopmentUnit resource tag and aws PrincipalTag/DevelopmentUnit
- C. Pass an attribute for DevelopmentUnit as an AWS Security Token Service (AWS STS) session tag during SAML federation Create an SCP with an allow action and a StrmgEquals condition for the DevelopmentUnit resource tag and aws Principal Tag 'DevelopmentUnit Assign the SCP to the root OU.
- D. Create separate IAM policies for each development unit For every IAM policy add an allow action and a StringEquals condition for the DevelopmentUnit resource tag and the development unit name During SAML federation use AWS Security Token Service (AWS STS) to assign the IAM policy and match the development unit name to the assumed IAM role

Answer: A

NEW QUESTION 106

- (Exam Topic 2)

A company is planning to migrate an application from on premises to AWS. The application currently uses an Oracle database and the company can tolerate a brief downtime of 1 hour when performing the switch to the new infrastructure As part of the migration. the database engine will be changed to MySQL. A solutions architect needs to determine which AWS services can be used to perform the migration while minimizing the amount of work and time required.

Which of the following will meet the requirements?

- A. Use AWS SCT to generate the schema scripts and apply them on the target prior to migration Use AWS DMS to analyse the current schema and provide a recommendation for the optimal database engine Then, use AWS DMS to migrate to the recommended engine Use AWS SCT to identify what embedded SQL code in the application can be converted and what has to be done manually
- B. Use AWS SCT to generate the schema scripts and apply them on the target prior to migratio
- C. Use AWS DMS to begin moving data from the on-premises database to AW
- D. After the initial copy continue to use AWS DMS to keep the databases m sync until cutting over to the new database Use AWS SCT to identify what embedded SQL code in the application can be converted and what has to be done manually.
- E. Use AWS DMS lo help identify the best target deployment between installing the database engine on Amazon EC2 directly or moving to Amazon RD
- F. Then, use AWS DMS to migrate to the platfor
- G. Use AWS Application Discovery Service to identify what embedded SQL code in the application can be converted and what has to be done manually.
- H. Use AWS DMS to begin moving data from the on-premises database to AWS After the initial copy, continue to use AWS DMS to keep the databases in sync until cutting over to the new database use AWS Application Discovery Service to identify what embedded SQL code m the application can be converted and what has to be done manually

Answer: B

NEW QUESTION 110

- (Exam Topic 2)

A large company runs workloads in VPCs that are deployed across hundreds of AWS accounts Each VPC consists of public subnets and private subnets that span across multiple Availability Zones NAT gateways are deployed in the public subnets and allow outbound connectivity to the internet from the private subnets.

A solutions architect is working on a hub-and-spoke design. All private subnets in the spoke VPCs must route traffic to the internet through an egress VPC The solutions architect already has deployed a NAT gateway in an egress VPC in a central AWS account

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

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

Answer: B

NEW QUESTION 111

- (Exam Topic 2)

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

The company needs to improve its ability to support the additional usage while minimizing the increase in costs associated with the solution.

Which strategy meets these requirements?

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

Answer: A

NEW QUESTION 114

- (Exam Topic 2)

A company is migrating an application to the AWS Cloud. The application runs in an on-premises data center and writes thousands of images into a mounted NFS file system each night After the company migrates the application, the company will host the application on an Amazon EC2 instance with a mounted Amazon Elastic File System (Amazon EFS) file system.

The company has established an AWS Direct Connect connection to AWS Before the migration cutover, a solutions architect must build a process that will replicate the newly created on-premises images to the EFS file system

What is the MOST operationally efficient way to replicate the images?

- A. Configure a periodic process to run the aws s3 sync command from the on-premises file system to Amazon S3 Configure an AWS Lambda function to process event notifications from Amazon S3 and copy the images from Amazon S3 to the EFS file system
- B. Deploy an AWS Storage Gateway file gateway with an NFS mount point
- C. Mount the file gateway file system on the on-premises server
- D. Configure a process to periodically copy the images to the mount point
- E. Deploy an AWS DataSync agent to an on-premises server that has access to the NFS file system Send data over the Direct Connect connection to an S3 bucket by using a public VIF Configure an AWS Lambda function to process event notifications from Amazon S3 and copy the images from Amazon S3 to the EFS file system
- F. Deploy an AWS DataSync agent to an on-premises server that has access to the NFS file system Send data over the Direct Connect connection to an AWS PrivateLink interface VPC endpoint for Amazon EFS by using a private VIF Configure a DataSync scheduled task to send the images to the EFS file system every 24 hours.

Answer: A

NEW QUESTION 119

- (Exam Topic 2)

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

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

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

Answer: ADF

NEW QUESTION 122

- (Exam Topic 2)

A solutions architect must update an application environment within AWS Elastic Beanstalk using a With green deployment methodology. The solutions architect creates an environment that is identical to the existing application environment and deploys the application to the new environment.

What should be done next to complete the update?

- A. Redirect to the new environment using Amazon Route 53
- B. Select the Swap Environment URLs option.
- C. Replace the Auto Scaling launch configuration
- D. Update the DNS records to point to the green environment

Answer: B

NEW QUESTION 123

- (Exam Topic 2)

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

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

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

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

Answer: B

NEW QUESTION 126

- (Exam Topic 2)

A company runs its application in the eu-west-1 Region and has one account for each of its environments: development, testing, and production. All the environments are running 24 hours a day, 7 days a week by using stateful Amazon EC2 instances and Amazon RDS for MySQL databases. The databases are between 500 GB and 800 GB in size.

The development team and testing team work on business days during business hours, but the production environment operates 24 hours a day, 7 days a week.

The company wants to reduce costs. All resources are tagged with an environment tag with either development, testing, or production as the key.

What should a solutions architect do to reduce costs with the LEAST operational effort?

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs once every day. Configure the rule to invoke one AWS Lambda function that starts or stops instances based on the tag, day, and time.
- B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs every business day in the evening.
- C. Configure the rule to invoke an AWS Lambda function that stops instances based on the tag. Create a second EventBridge (CloudWatch Events) rule that runs every business day in the morning. Configure the second rule to invoke another Lambda function that starts instances based on the tag.
- D. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs every business day in the evening. Configure the rule to invoke an AWS Lambda function that terminates instances based on the tag. Create a second EventBridge (CloudWatch Events) rule that runs every business day in the morning. Configure the second rule to invoke another Lambda function that restores the instances from their last backup based on the tag.
- E. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs every hour. Configure the rule to invoke one AWS Lambda function that terminates or restores instances from their last backup based on the tag, day, and time.
- F. day, and time

Answer: C

NEW QUESTION 131

- (Exam Topic 2)

A company is hosting a three-tier web application in an on-premises environment. Due to a recent surge in traffic that resulted in downtime and a significant financial impact, company management has ordered that the application be moved to AWS. The application is written in .NET and has a dependency on a MySQL database. A solutions architect must design a scalable and highly available solution to meet the demand of 200,000 daily users.

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

- A. Use AWS Elastic Beanstalk to create a new application with a web server environment and an Amazon RDS MySQL Multi-AZ DB instance. The environment should launch a Network Load Balancer (NLB) in front of an Amazon EC2 Auto Scaling group in multiple Availability Zones. Use an Amazon Route 53 alias record to route traffic from the company's domain to the NLB.
- B. Use AWS CloudFormation to launch a stack containing an Application Load Balancer (ALB) in front of an Amazon EC2 Auto Scaling group spanning three Availability Zones.
- C. The stack should launch a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a Retain deletion policy.
- D. Use an Amazon Route 53 alias record to route traffic from the company's domain to the ALB.
- E. Use AWS Elastic Beanstalk to create an automatically scaling web server environment that spans two separate Regions with an Application Load Balancer (ALB) in each Region.
- F. Create a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a cross-Region read replica. Use Amazon Route 53 with a geoproximity routing policy to route traffic between the two Regions.
- G. Use AWS CloudFormation to launch a stack containing an Application Load Balancer (ALB) in front of an Amazon ECS cluster of Spot Instances spanning three Availability Zones. The stack should launch an Amazon RDS MySQL DB instance with a Snapshot deletion policy. Use an Amazon Route 53 alias record to route traffic from the company's domain to the ALB.

Answer: B

NEW QUESTION 135

- (Exam Topic 2)

A company deploys a new web application. As part of the setup, the company configures AWS WAF to log to Amazon S3 through Amazon Kinesis Data Firehose. The company develops an Amazon Athena query that runs once daily to return AWS WAF log data from the previous 24 hours. The volume of daily logs is constant. However, over time, the same query is taking more time to run.

A solutions architect needs to design a solution to prevent the query time from continuing to increase. The solution must minimize operational overhead.

Which solution will meet these requirements?

- A. Create an AWS Lambda function that consolidates each day's AWS WAF logs into one log file
- B. Reduce the amount of data scanned by configuring AWS WAF to send logs to a different S3 bucket each day
- C. Update the Kinesis Data Firehose configuration to partition the data in Amazon S3 by date and time Create external tables for Amazon Redshift Configure Amazon Redshift Spectrum to query the data source
- D. Modify the Kinesis Data Firehose configuration and Athena table definition to partition the data by date and time
- E. Change the Athena query to view the relevant partitions

Answer: D

NEW QUESTION 136

- (Exam Topic 2)

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

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

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

Answer: B

NEW QUESTION 137

- (Exam Topic 2)

A retail company runs a business-critical web service on an Amazon Elastic Container Service (Amazon ECS) cluster that runs on Amazon EC2 instances. The web service receives POST requests from end users and writes data to a MySQL database that runs on a separate EC2 instance. The company needs to ensure that data loss does not occur.

The current code deployment process includes manual updates of the ECS service. During a recent deployment, end users encountered intermittent 502 Bad Gateway errors in response to valid web requests.

The company wants to implement a reliable solution to prevent this issue from recurring. The company also wants to automate code deployments. The solution must be highly available and must optimize cost-effectiveness.

- A. Run the web service on an ECS cluster that has a Fargate launch type. Use AWS CodePipeline and AWS CodeDeploy to perform a blue/green deployment with validation testing to update the ECS service.
- B. Migrate the MySQL database to run on an Amazon RDS for MySQL Multi-AZ DB instance that uses Provisioned IOPS SSD (io2) storage.
- C. Configure an Amazon Simple Queue Service (Amazon SQS) queue as an event source to receive the POST requests from the web service. Configure an AWS Lambda function to poll the queue. Write the data to the database.
- D. Run the web service on an ECS cluster that has a Fargate launch type. Use AWS CodePipeline and AWS CodeDeploy to perform a canary deployment to update the ECS service.

Answer: CD

NEW QUESTION 142

- (Exam Topic 2)

A company runs an application in the cloud that consists of a database and a website. Users can post data to the website, have the data processed, and have the data sent back to them in an email. Data is stored in a MySQL database running on an Amazon EC2 instance. The database is running in a VPC with two private subnets. The website is running on Apache Tomcat in a single EC2 instance in a different VPC with one public subnet. There is a single VPC peering connection between the database and website VPC.

The website has suffered several outages during the last month due to high traffic.

Which actions should a solutions architect take to increase the reliability of the application? (Select THREE.)

- A. Place the Tomcat server in an Auto Scaling group with multiple EC2 instances behind an Application Load Balancer
- B. Provision an additional VPC peering connection
- C. Migrate the MySQL database to Amazon Aurora with one Aurora Replica
- D. Provision two NAT gateways in the database VPC
- E. Move the Tomcat server to the database VPC
- F. Create an additional public subnet in a different Availability Zone in the website VPC

Answer: ACF

NEW QUESTION 146

- (Exam Topic 2)

A company plans to deploy a new private Intranet service on Amazon EC2 instances inside a VPC. An AWS Site-to-Site VPN connects the VPC to the company's on-premises network. The new service must communicate with existing on-premises services. The on-premises services are accessible through the use of hostnames that reside in the company's example DNS zone. This DNS zone is wholly hosted on premises and is available only on the company's private network. A solutions architect must ensure that the new service can resolve hostnames on the company.example domain to integrate with existing services. Which solution meets these requirements?

- A. Create an empty private zone in Amazon Route 53 for company.example
- B. Add an additional NS record to the company's on-premises company.example zone that points to the authoritative name servers for the new private zone in Route 53
- C. Turn on DNS hostnames for the VPC
- D. Configure a new outbound endpoint with Amazon Route 53 Resolver
- E. Create a Resolver rule to forward requests for company.example to the on-premises name servers.
- F. Turn on DNS hostnames for the VPC
- G. Configure a new inbound resolver endpoint with Amazon Route 53 Resolver
- H. Configure the on-premises DNS server to forward requests for company.example to the new resolver.

- I. Use AWS Systems Manager to configure a run document that will install a hosts file that contains any required hostname
J. Use an Amazon EventBridge (Amazon CloudWatch Events) rule to run the document when an instance is entering the running state.

Answer: C

NEW QUESTION 151

- (Exam Topic 2)

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

A solutions architect designs a solution in which the IoT devices are streaming to Amazon Kinesis Data Streams. Several applications are reading from the stream. However, several consumers are experiencing throttling and are periodically encountering a ReadProvisionedThroughputExceeded error.

Which actions should the solutions architect take to resolve this issue? (Select THREE.)

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

Answer: ACD

NEW QUESTION 152

- (Exam Topic 2)

A company has used infrastructure as code (IaC) to provision a set of two Amazon EC2 instances. The instances have remained the same for several years.

The company's business has grown rapidly in the past few months. In response, the company's operations team has implemented an Auto Scaling group to manage the sudden increases in traffic. Company policy requires a monthly installation of security updates on all operating systems that are running.

The most recent security update required a reboot. As a result the Auto Scaling group terminated the instances and replaced them with new, unpatched instances. Which combination of steps should a solutions architect recommend to avoid a recurrence of this issue? (Select TWO)

- A. Modify the Auto Scaling group by setting the Update policy to target the oldest launch configuration for replacement.
- B. Create a new Auto Scaling group before the next patch maintenance. During the maintenance window patch both groups and reboot the instances.
- C. Create an Elastic Load Balancer in front of the Auto Scaling group. Configure monitoring to ensure that target group health checks return healthy after the Auto Scaling group replaces the terminated instances.
- D. Create automation scripts to patch an AMI.
- E. Update the launch configuration, and invoke an Auto Scaling instance refresh.
- F. Create an Elastic Load Balancer in front of the Auto Scaling group. Configure termination protection on the instances.

Answer: AC

NEW QUESTION 153

- (Exam Topic 2)

A solutions architect has implemented a SAML 2.0 federated identity solution with their company's

on-premises identity provider (IdP) to authenticate users' access to the AWS environment. When the solutions architect tests authentication through the federated identity web portal, access to the AWS environment is granted. However, when test users attempt to authenticate through the federated identity web portal, they are not able to access the AWS environment.

Which items should the solutions architect check to ensure identity federation is properly configured? (Select THREE)

- A. The IAM user's permissions policy has allowed the use of SAML federation for that user.
- B. The IAM roles created for the federated users' or federated groups' trust policy have set the SAML provider as the principal.
- C. Test users are not in the AWSFederatedUsers group in the company's IdP.
- D. The web portal calls the AWS STS AssumeRoleWithSAML API with the ARN of the SAML provider, the ARN of the IAM role, and the SAML assertion from IdP.
- E. The on-premises IdP's DNS hostname is reachable from the AWS environment VPCs.
- F. The company's IdP defines SAML assertions that properly map users or groups in the company to IAM roles with appropriate permissions.

Answer: BCF

NEW QUESTION 157

- (Exam Topic 2)

A company is creating a sequel for a popular online game. A large number of users from all over the world will play the game within the first week after launch.

Currently, the game consists of the following components deployed in a single AWS Region:

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

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

What should the solutions architect do to meet these requirements?

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

Answer: C

NEW QUESTION 161

- (Exam Topic 2)

A data analytics company has an Amazon Redshift cluster that consists of several reserved nodes. The cluster is experiencing unexpected bursts of usage because a team of employees is compiling a deep audit analysis report. The queries to generate the report are complex read queries and are CPU intensive. Business requirements dictate that the cluster must be able to service read and write queries at any times. A solutions architect must devise a solution that accommodates the bursts of usage.

Which solution meets these requirements MOST cost-effectively?

- A. Provision an Amazon EMR cluster. Offload the complex data processing tasks.
- B. Deploy an AWS Lambda function to add capacity to the Amazon Redshift cluster by using a classic resize operation when the cluster's CPU metrics in Amazon CloudWatch reach 80%.
- C. Deploy an AWS Lambda function to add capacity to the Amazon Redshift cluster by using an elastic resize operation when the cluster's CPU metrics in Amazon CloudWatch reach 80%.
- D. Turn on the Concurrency Scaling feature for the Amazon Redshift cluster.

Answer: D

NEW QUESTION 166

- (Exam Topic 2)

A media storage application uploads user photos to Amazon S3 for processing by AWS Lambda functions. Application state is stored in Amazon DynamoDB tables. Users are reporting that some uploaded photos are not being processed properly. The application developers trace the logs and find that Lambda is experiencing photo processing issues when thousands of users upload photos simultaneously. The issues are the result of Lambda concurrency limits and the performance of DynamoDB when data is saved.

Which combination of actions should a solutions architect take to increase the performance and reliability of the application? (Select TWO.)

- A. Evaluate and adjust the RCUs for the DynamoDB tables.
- B. Evaluate and adjust the WCUs for the DynamoDB tables.
- C. Add an Amazon ElastiCache layer to increase the performance of Lambda functions.
- D. Add an Amazon Simple Queue Service (Amazon SQS) queue and reprocessing logic between Amazon S3 and the Lambda functions.
- E. Use S3 Transfer Acceleration to provide lower latency to users.

Answer: BD

NEW QUESTION 168

- (Exam Topic 2)

A company is finalizing the architecture for its backup solution for applications running on AWS. All of the applications run on AWS and use at least two Availability Zones in each tier.

Company policy requires IT to durably store nightly backups of all its data in at least two locations: production and disaster recovery. The locations must be in different geographic regions. The company also needs the backup to be available to restore immediately at the production data center, and within 24 hours at the disaster recovery location. Backup processes must be fully automated.

What is the MOST cost-effective backup solution that will meet all requirements?

- A. Back up all the data to a large Amazon EBS volume attached to the backup media server in the production region.
- B. Run automated scripts to snapshot these volumes nightly.
- C. and copy these snapshots to the disaster recovery region.
- D. Back up all the data to Amazon S3 in the disaster recovery region. Use a Lifecycle policy to move this data to Amazon Glacier in the production region immediately. Only the data is replicated; remove the data from the S3 bucket in the disaster recovery region.
- E. Back up all the data to Amazon Glacier in the production region.
- F. Set up cross-region replication of this data to Amazon Glacier in the disaster recovery region.
- G. Set up a lifecycle policy to delete any data older than 60 days.
- H. Back up all the data to Amazon S3 in the production region.
- I. Set up cross-region replication of this S3 bucket to another region and set up a lifecycle policy in the second region to immediately move this data to Amazon Glacier.

Answer: D

NEW QUESTION 173

- (Exam Topic 2)

A company is migrating an on-premises content management system (CMS) to AWS Fargate. The company uses the CMS for blog posts that include text, images, and videos. The company has observed that traffic to blog posts drops by more than 80% after the posts are more than 30 days old.

The CMS runs on multiple VMs and stores application state on disk. This application state is shared across all instances across multiple Availability Zones. Images and other media are stored on a separate NFS file share. The company needs to reduce the costs of the existing solution while minimizing the impact on performance.

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

- A. Store media in an Amazon S3 Standard bucket. Create an S3 Lifecycle configuration that transitions objects that are older than 30 days to the S3 Standard-Infrequent Access (S3 Standard-IA) storage class.
- B. Store media on an Amazon Elastic File System (Amazon EFS) volume. Attach the EFS volume to all Fargate instances.
- C. Store application state on an Amazon Elastic File System (Amazon EFS) volume. Attach the EFS volume to all Fargate instances.
- D. Store application state on an Amazon Elastic Block Store (Amazon EBS) volume. Attach the EBS volume to all Fargate instances.
- E. Store media in an Amazon S3 Standard bucket. Create an S3 Lifecycle configuration that transitions objects that are older than 30 days to the S3 Glacier storage class.

Answer: AC

NEW QUESTION 174

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your SAP-C02 Exam with Our Prep Materials Via below:

<https://www.certleader.com/SAP-C02-dumps.html>