

Amazon-Web-Services

Exam Questions SAP-C02

AWS Certified Solutions Architect - Professional



NEW QUESTION 1

- (Exam Topic 1)

A startup company recently migrated a large ecommerce website to AWS. The website has experienced a 70% increase in sales. Software engineers are using a private GitHub repository to manage code. The DevOps learn is using Jenkins for builds and unit testing. The engineers need to receive notifications for bad builds and zero downtime during deployments. The engineers also need to ensure any changes to production are seamless for users and can be rolled back in the event of a major issue.

The software engineers have decided to use AWS CodePipeline to manage their build and deployment process.

Which solution will meet these requirements?

- A. Use GitHub websockets to trigger the CodePipeline pipeline
- B. Use the Jenkins plugin for AWS CodeBuild to conduct unit testing
- C. Send alerts to an Amazon SNS topic for any bad build
- D. Deploy in an in-place
- E. all-at-once deployment configuration using AWS CodeDeploy.
- F. Use GitHub webhooks to trigger the CodePipeline pipeline
- G. Use the Jenkins plugin for AWS CodeBuild to conduct unit testing
- H. Send alerts to an Amazon SNS topic for any bad build
- I. Deploy in a blue/green deployment using AWS CodeDeploy.
- J. Use GitHub websockets to trigger the CodePipeline pipeline
- K. Use AWS X-Ray for unit testing and static code analysis
- L. Send alerts to an Amazon SNS topic for any bad build
- M. Deploy in a blue/green deployment using AWS CodeDeploy.
- N. Use GitHub webhooks to trigger the CodePipeline pipeline
- O. Use AWS X-Ray for unit testing and static code analysis
- P. Send alerts to an Amazon SNS topic for any bad build
- Q. Deploy in an in-place, all-at-once deployment configuration using AWS CodeDeploy.

Answer: B

NEW QUESTION 2

- (Exam Topic 1)

A company is running an application distributed over several Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer. The security team requires that all application access attempts be made available for analysis. Information about the client IP address, connection type, and user agent must be included.

Which solution will meet these requirements?

- A. Enable EC2 detailed monitoring, and include network logs. Send all logs through Amazon Kinesis Data Firehose to an Amazon Elasticsearch Service (Amazon ES) cluster that the security team uses for analysis.
- B. Enable VPC Flow Logs for all EC2 instance network interfaces. Publish VPC Flow Logs to an Amazon S3 bucket. Have the security team use Amazon Athena to query and analyze the logs.
- C. Enable access logs for the Application Load Balancer, and publish the logs to an Amazon S3 bucket. Have the security team use Amazon Athena to query and analyze the logs.
- D. Enable Traffic Mirroring and specify all EC2 instance network interfaces as the source.
- E. Send all traffic information through Amazon Kinesis Data Firehose to an Amazon Elasticsearch Service (Amazon ES) cluster that the security team uses for analysis.

Answer: C

Explanation:

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-access-logs.html>

NEW QUESTION 3

- (Exam Topic 1)

A company has a complex web application that leverages Amazon CloudFront for global scalability and performance. Over time, users report that the web application is slowing down.

The company's operations team reports that the CloudFront cache hit ratio has been dropping steadily. The cache metrics report indicates that query strings on some URLs are inconsistently ordered and are specified sometimes in mixed-case letters and sometimes in lowercase letters.

Which set of actions should the solutions architect take to increase the cache hit ratio as quickly as possible?

- A. Deploy a Lambda@Edge function to sort parameters by name and force them to be lowercase.
- B. Select the CloudFront viewer request trigger to invoke the function.
- C. Update the CloudFront distribution to disable caching based on query string parameters.
- D. Deploy a reverse proxy after the load balancer to post-process the emitted URLs in the application to force the URL strings to be lowercase.
- E. Update the CloudFront distribution to specify casing-insensitive query string processing.

Answer: A

Explanation:

https://docs.amazonaws.cn/en_us/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html#lambda-ex Before CloudFront serves content from the cache it will trigger any Lambda function associated with the Viewer Request, in which we can normalize parameters.

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html#lambda-examp>

NEW QUESTION 4

- (Exam Topic 1)

A company has an Amazon VPC that is divided into a public subnet and a private subnet. A web application runs in Amazon VPC, and each subnet has its own NACL. The public subnet has a CIDR of 10.0.0.0/24. An Application Load Balancer is deployed to the public subnet. The private subnet has a CIDR of 10.0.1.0/24. Amazon EC2 instances that run a web server on port 80 are launched into the private subnet.

Only network traffic that is required for the Application Load Balancer to access the web application can be allowed to travel between the public and private

subnets

What collection of rules should be written to ensure that the private subnet's NACL meets the requirement? (Select TWO.)

- A. An inbound rule for port 80 from source 0.0.0.0/0
- B. An inbound rule for port 80 from source 10.0.0.0/24
- C. An outbound rule for port 80 to destination 0.0.0.0/0
- D. An outbound rule for port 80 to destination 10.0.0.0/24
- E. An outbound rule for ports 1024 through 65535 to destination 10.0.0.0/24

Answer: BE

Explanation:

Ephemeral ports are not covered in the syllabus so be careful that you don't confuse day to day best practise with what is required for the exam. Link to an explanation on Ephemeral ports here. <https://acloud.guru/forums/aws-certified-solutions-architect-associate/discussion/-KUbcwo4lXefMI7janaK/netw>

NEW QUESTION 5

- (Exam Topic 1)

A company is running a web application on Amazon EC2 instances in a production AWS account. The company requires all logs generated from the web application to be copied to a central AWS account (or analysis and archiving). The company's AWS accounts are currently managed independently. Logging agents are configured on the EC2 instances to upload the log files to an Amazon S3 bucket in the central AWS account.

A solutions architect needs to provide access for a solution that will allow the production account to store log files in the central account. The central account also needs to have read access to the log files.

What should the solutions architect do to meet these requirements?

- A. Create a cross-account role in the central account
- B. Assume the role from the production account when the logs are being copied.
- C. Create a policy on the S3 bucket with the production account ID as the principal
- D. Allow S3 access from a delegated user.
- E. Create a policy on the S3 bucket with access from only the CIDR range of the EC2 instances in the production account
- F. Use the production account ID as the principal.
- G. Create a cross-account role in the production account
- H. Assume the role from the production account when the logs are being copied.

Answer: B

NEW QUESTION 6

- (Exam Topic 1)

A company is running a data-intensive application on AWS. The application runs on a cluster of hundreds of Amazon EC2 instances. A shared file system also runs on several EC2 instances that store 200 TB of data. The application reads and modifies the data on the shared file system and generates a report. The job runs once monthly, reads a subset of the files from the shared file system, and takes about 72 hours to complete. The compute instances scale in an Auto Scaling group, but the instances that host the shared file system run continuously. The compute and storage instances are all in the same AWS Region.

A solutions architect needs to reduce costs by replacing the shared file system instances. The file system must provide high performance access to the needed data for the duration of the 72-hour run.

Which solution will provide the LARGEST overall cost reduction while meeting these requirements?

- A. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Intelligent-Tiering storage class
- B. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using lazy loadin
- C. Use the new file system as the shared storage for the duration of the job
- D. Delete the file system when the job is complete.
- E. Migrate the data from the existing shared file system to a large Amazon Elastic Block Store (Amazon EBS) volume with Multi-Attach enable
- F. Attach the EBS volume to each of the instances by using a user data script in the Auto Scaling group launch template
- G. Use the EBS volume as the shared storage for the duration of the job
- H. Detach the EBS volume when the job is complete.
- I. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Standard storage class
- J. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using batch loadin
- K. Use the new file system as the shared storage for the duration of the job
- L. Delete the file system when the job is complete.
- M. Migrate the data from the existing shared file system to an Amazon S3 bucket
- N. Before the job runs each month, use AWS Storage Gateway to create a file gateway with the data from Amazon S3. Use the file gateway as the shared storage for the job
- O. Delete the file gateway when the job is complete.

Answer: B

NEW QUESTION 7

- (Exam Topic 1)

A company is developing and hosting several projects in the AWS Cloud. The projects are developed across multiple AWS accounts under the same organization in AWS Organizations. The company requires the cost for cloud infrastructure to be allocated to the owning project. The team responsible for all of the AWS accounts has discovered that several Amazon EC2 instances are lacking the Project tag used for cost allocation.

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

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

Answer: BDE

NEW QUESTION 8

- (Exam Topic 1)

A company built an ecommerce website on AWS using a three-tier web architecture. The application is Java-based and composed of an Amazon CloudFront distribution, an Apache web server layer of Amazon EC2 instances in an Auto Scaling group, and a backend Amazon Aurora MySQL database.

Last month, during a promotional sales event, users reported errors and timeouts while adding items to their shopping carts. The operations team recovered the logs created by the web servers and reviewed Aurora DB cluster performance metrics. Some of the web servers were terminated before logs could be collected and the Aurora metrics were not sufficient for query performance analysis.

Which combination of steps must the solutions architect take to improve application performance visibility during peak traffic events? (Select THREE.)

- A. Configure the Aurora MySQL DB cluster to publish slow query and error logs to Amazon CloudWatch Logs.
- B. Implement the AWS X-Ray SDK to trace incoming HTTP requests on the EC2 instances and implement tracing of SQL queries with the X-Ray SDK for Java.
- C. Configure the Aurora MySQL DB cluster to stream slow query and error logs to Amazon Kinesis.
- D. Install and configure an Amazon CloudWatch Logs agent on the EC2 instances to send the Apache logs to CloudWatch Logs.
- E. Enable and configure AWS CloudTrail to collect and analyze application activity from Amazon EC2 and Aurora.
- F. Enable Aurora MySQL DB cluster performance benchmarking and publish the stream to AWS X-Ray.

Answer: ABD

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_LogAccess.Concepts.MySQL.html <https://aws.amazon.com/blogs/mt/simplifying-apache-server-logs-with-amazon-cloudwatch-logs-insights/> <https://docs.aws.amazon.com/xray/latest/devguide/xray-sdk-dotnet-messagehandler.html>
<https://docs.aws.amazon.com/xray/latest/devguide/xray-sdk-java-sqlclients.html>

NEW QUESTION 9

- (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 10

- (Exam Topic 1)

A solutions architect is responsible for redesigning a legacy Java application to improve its availability, data durability, and scalability. Currently, the application runs on a single high-memory Amazon EC2 instance. It accepts HTTP requests from upstream clients, adds them to an in-memory queue, and responds with a 200 status. A separate application thread reads items from the queue, processes them, and persists the results to an Amazon RDS MySQL instance. The processing time for each item takes 90 seconds on average, most of which is spent waiting on external service calls, but the application is written to process multiple items in parallel.

Traffic to this service is unpredictable. During periods of high load, items may sit in the internal queue for over an hour while the application processes the backlog. In addition, the current system has issues with availability and data loss if the single application node fails.

Clients that access this service cannot be modified. They expect to receive a response to each HTTP request they send within 10 seconds before they will time out and retry the request.

Which approach would improve the availability and durability of the system while decreasing the processing latency and minimizing costs?

- A. Create an Amazon API Gateway REST API that uses Lambda proxy integration to pass requests to an AWS Lambda function.
- B. Migrate the core processing code to a Lambda function and write a wrapper class that provides a handler method that converts the proxy events to the internal application data model and invokes the processing module.
- C. Create an Amazon API Gateway REST API that uses a service proxy to put items in an Amazon SQS queue.
- D. Extract the core processing code from the existing application and update it to pull items from Amazon SQS instead of an in-memory queue.
- E. Deploy the new processing application to smaller EC2 instances within an Auto Scaling group that scales dynamically based on the approximate number of messages in the Amazon SQS queue.
- F. Modify the application to use Amazon DynamoDB instead of Amazon RDS.
- G. Configure Auto Scaling for the DynamoDB table.
- H. Deploy the application within an Auto Scaling group with a scaling policy based on CPU utilization.
- I. Back the in-memory queue with a memory-mapped file to an instance store volume and periodically write that file to Amazon S3.
- J. Update the application to use a Redis task queue instead of the in-memory queue.
- K. Build a Docker container image for the application.
- L. Create an Amazon ECS task definition that includes the application container and a separate container to host Redis.
- M. Deploy the new task definition as an ECS service using AWS Fargate, and enable Auto Scaling.

Answer: B

Explanation:

The obvious challenges here are long workloads, scalability based on queue load, and reliability. Almost always the de facto answer to queue-related workload is SQS. Since the workloads are very long (90 minutes) Lambdas cannot be used (15 mins max timeout). So, autoscaled smaller EC2 nodes that wait on external

services to complete the task makes more sense. If the task fails, the message is returned to the queue and retried.

NEW QUESTION 10

- (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 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
- 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 11

- (Exam Topic 1)

A large company is running a popular web application. The application runs on several Amazon EC2 Linux Instances in an Auto Scaling group in a private subnet. An Application Load Balancer is targeting the Instances in the Auto Scaling group in the private subnet. AWS Systems Manager Session Manager is configured, and AWS Systems Manager Agent is running on all the EC2 instances.

The company recently released a new version of the application. Some EC2 instances are now being marked as unhealthy and are being terminated. As a result, the application is running at reduced capacity. A solutions architect tries to determine the root cause by analyzing Amazon CloudWatch logs that are collected from the application, but the logs are inconclusive.

How should the solutions architect gain access to an EC2 instance to troubleshoot the issue?

- A. Suspend the Auto Scaling group's HealthCheck scaling process
- B. Use Session Manager to log in to an instance that is marked as unhealthy
- C. Enable EC2 instance termination protection. Use Session Manager to log in to an instance that is marked as unhealthy.
- D. Set the termination policy to OldestInstance on the Auto Scaling group
- E. Use Session Manager to log in to an instance that is marked as unhealthy
- F. Suspend the Auto Scaling group's Terminate process
- G. Use Session Manager to log in to an instance that is marked as unhealthy

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 12

- (Exam Topic 1)

A company runs an application that gives users the ability to search for videos and related information by using keywords that are curated from content providers. The application data is stored in an on-premises Oracle database that is 800 GB in size.

The company wants to migrate the data to an Amazon Aurora MySQL DB instance. A solutions architect plans to use the AWS Schema Conversion Tool and AWS Database Migration Service (AWS DMS) for the migration. During the migration, the existing database must serve ongoing requests. The migration must be completed with minimum downtime.

Which solution will meet these requirements?

- A. Create primary key indexes, secondary indexes, and referential integrity constraints in the target database before starting the migration process
- B. Use AWS DMS to run the conversion report for Oracle to Aurora MySQL
- C. Remediate any issues. Then use AWS DMS to migrate the data
- D. Use the M5 or CS DMS replication instance type for ongoing replication
- E. Turn off automatic backups and logging of the target database until the migration and cutover processes are complete

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Managing.Backups.html>

NEW QUESTION 14

- (Exam Topic 1)

A company provides a centralized Amazon EC2 application hosted in a single shared VPC. The centralized application must be accessible from client applications running in the VPCs of other business units. The centralized application front end is configured with a Network Load Balancer (NLB) for scalability. Up to 10 business unit VPCs will need to be connected to the shared VPC. Some of the business unit VPC CIDR blocks overlap with the shared VPC, and some overlap with each other. Network connectivity to the centralized application in the shared VPC should be allowed from authorized business unit VPCs only. Which network configuration should a solutions architect use to provide connectivity from the client applications in the business unit VPCs to the centralized application in the shared VPC?

- A. Create an AWS Transit Gateway
- B. Attach the shared VPC and the authorized business unit VPCs to the transit gateway
- C. Create a single transit gateway route table and associate it with all of the attached VPC
- D. Allow automatic propagation of routes from the attachments into the route table
- E. Configure VPC routing tables to send traffic to the transit gateway.
- F. Create a VPC endpoint service using the centralized application NLB and enable the option to require endpoint acceptance
- G. Create a VPC endpoint in each of the business unit VPCs using the service name of the endpoint service
- H. Accept authorized endpoint requests from the endpoint service console.
- I. Create a VPC peering connection from each business unit VPC to the shared VPC
- J. Accept the VPC peering connections from the shared VPC console
- K. Configure VPC routing tables to send traffic to the VPC peering connection.
- L. Configure a virtual private gateway for the shared VPC and create customer gateways for each of the authorized business unit VPC
- M. Establish a Site-to-Site VPN connection from the business unit VPCs to the shared VPC
- N. Configure VPC routing tables to send traffic to the VPN connection.

Answer: B

Explanation:

Amazon Transit Gateway doesn't support routing between Amazon VPCs with overlapping CIDRs. If you attach a new Amazon VPC that has a CIDR which overlaps with an already attached Amazon VPC, Amazon Transit Gateway will not propagate the new Amazon VPC route into the Amazon Transit Gateway route table.

<https://docs.aws.amazon.com/elasticloadbalancing/latest/network/load-balancer-target-groups.html#client-ip-pre>

NEW QUESTION 19

- (Exam Topic 1)

A company is migrating an application to AWS. It wants to use fully managed services as much as possible during the migration. The company needs to store large, important documents within the application with the following requirements:

- * 1. The data must be highly durable and available.
- * 2. The data must always be encrypted at rest and in transit.
- * 3. The encryption key must be managed by the company and rotated periodically.

Which of the following solutions should the solutions architect recommend?

- A. Deploy the storage gateway to AWS in file gateway mode
- B. Use Amazon EBS volume encryption using an AWS KMS key to encrypt the storage gateway volumes.
- C. Use Amazon S3 with a bucket policy to enforce HTTPS for connections to the bucket and to enforce server-side encryption and AWS KMS for object encryption.
- D. Use Amazon DynamoDB with SSL to connect to DynamoDB
- E. Use an AWS KMS key to encrypt DynamoDB objects at rest.
- F. Deploy instances with Amazon EBS volumes attached to store this data
- G. Use EBS volume encryption using an AWS KMS key to encrypt the data.

Answer: B

Explanation:

Use Amazon S3 with a bucket policy to enforce HTTPS for connections to the bucket and to enforce server-side encryption and AWS KMS for object encryption.

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 replica
- 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 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 replica. Set the alarm to directly invoke an AWS Lambda function to delete its Route 53 record set.
- F. 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 27

- (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 configuration
- 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 configuration
- 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 nodes
- 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 32

- (Exam Topic 1)

A company's AWS architecture currently uses access keys and secret access keys stored on each instance to access AWS services. Database credentials are hard-coded on each instance. SSH keys for command-line remote access are stored in a secured Amazon S3 bucket. The company has asked its solutions architect to improve the security posture of the architecture without adding operational complexity.

Which combination of steps should the solutions architect take to accomplish this? (Select THREE.)

- A. Use Amazon EC2 instance profiles with an IAM role.
- B. Use AWS Secrets Manager to store access keys and secret access keys.
- C. Use AWS Systems Manager Parameter Store to store database credentials.
- D. Use a secure fleet of Amazon EC2 bastion hosts (or remote access).
- E. Use AWS KMS to store database credentials.
- F. Use AWS Systems Manager Session Manager for remote access

Answer: ACF

Explanation:

<https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager.html>

NEW QUESTION 33

- (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 36

- (Exam Topic 1)

A start up company hosts a fleet of Amazon EC2 instances in private subnets using the latest Amazon Linux 2 AMI. The company's engineers rely heavily on SSH access to the instances for troubleshooting.

The company's existing architecture includes the following:

- A VPC with private and public subnets, and a NAT gateway
- Site-to-Site VPN for connectivity with the on-premises environment
- EC2 security groups with direct SSH access from the on-premises environment

The company needs to increase security controls around SSH access and provide auditing of commands executed by the engineers.

Which strategy should a solutions architect use?

- A. Install and configure EC2 Instance Connect on the fleet of EC2 instance
- B. Remove all security group rules attached to EC2 instances that allow inbound TCP on port 22. Advise the engineers to remotely access the instances by using the EC2 Instance Connect CLI.
- C. Update the EC2 security groups to only allow inbound TCP on port 22 to the IP addresses of the engineer's device
- D. Install the Amazon CloudWatch agent on all EC2 instances and send operating system audit logs to CloudWatch Logs.
- E. Update the EC2 security groups to only allow inbound TCP on port 22 to the IP addresses of the engineer's device
- F. Enable AWS Config for EC2 security group resource change
- G. Enable AWS Firewall Manager and apply a security group policy that automatically remediates changes to rules.
- H. Create an IAM role with the AmazonSSMManagedInstanceCore managed policy attached
- I. Attach the IAM role to all the EC2 instance
- J. Remove all security group rules attached to the EC2
- K. instances that allow inbound TCP on port 22. Have the engineers install the AWS Systems Manager Session Manager plugin for their devices and remotely access the instances by using the start-session API call from Systems Manager.

Answer: B

NEW QUESTION 40

- (Exam Topic 1)

A finance company hosts a data lake in Amazon S3. The company receives financial data records over SFTP each night from several third parties. The company runs its own SFTP server on an Amazon EC2 instance in a public subnet of a VPC. After the files are uploaded, they are moved to the data lake by a cron job that runs on the same instance. The SFTP server is reachable on DNS sftp.examWe.com through the use of Amazon Route 53.

What should a solutions architect do to improve the reliability and scalability of the SFTP solution?

- A. Move the EC2 instance into an Auto Scaling group
- B. Place the EC2 instance behind an Application Load Balancer (ALB). Update the DNS record sftp.example.com in Route 53 to point to the ALB.
- C. Migrate the SFTP server to AWS Transfer for SFT
- D. Update the DNS record sftp.example.com in Route 53 to point to the server endpoint hostname.
- E. Migrate the SFTP server to a file gateway in AWS Storage Gateway
- F. Update the DNS record sftp.example.com in Route 53 to point to the file gateway endpoint.
- G. Place the EC2 instance behind a Network Load Balancer (NLB). Update the DNS record sftp.example.com in Route 53 to point to the NLB.

Answer: B

NEW QUESTION 45

- (Exam Topic 1)

A company is hosting a single-page web application in the AWS Cloud. The company is using Amazon CloudFront to reach its goal audience. The CloudFront distribution has an Amazon S3 bucket that is configured as its origin. The static files for the web application are stored in this S3 bucket.

The company has used a simple routing policy to configure an Amazon Route 53 A record The record points to the CloudFront distribution The company wants to use a canary deployment release strategy for new versions of the application.

What should a solutions architect recommend to meet these requirements?

- A. Create a second CloudFront distribution for the new version of the application
- B. Update the Route 53 record to use a weighted routing policy.
- C. Create a Lambda@Edge function
- D. Configure the function to implement a weighting algorithm and rewrite the URL to direct users to a new version of the application.
- E. Create a second S3 bucket and a second CloudFront origin for the new S3 bucket Create a CloudFrontorigin group that contains both origins Configure origin weighting for the origin group.
- F. Create two Lambda@Edge functions
- G. Use each function to serve one of the application versions Set up a CloudFront weighted Lambda@Edge invocation policy

Answer: A

NEW QUESTION 49

- (Exam Topic 2)

A life sciences company is using a combination of open source tools to manage data analysis workflows and Docker containers running on servers in its on-premises data center to process genomics data. Sequencing data is generated and stored on a local storage area network (SAN), and then the data is processed. The research and development teams are running into capacity issues and have decided to re-architect their genomics analysis platform on AWS to scale based on workload demands and reduce the turnaround time from weeks to days. The company has a high-speed AWS Direct Connect connection. Sequencers will generate around 200 GB of data for each genome, and individual jobs can take several hours to process the data with ideal compute capacity. The end result will be stored in Amazon S3. The company is expecting 10-15 job requests each day. Which solution meets these requirements?

- A. Use regularly scheduled AWS Snowball Edge devices to transfer the sequencing data into AWS. When AWS receives the Snowball Edge device and the data is loaded into Amazon S3, use S3 events to trigger an AWS Lambda function to process the data.
- B. Use AWS Data Pipeline to transfer the sequencing data to Amazon S3. Use S3 events to trigger an Amazon EC2 Auto Scaling group to launch custom-AMI EC2 instances running the Docker containers to process the data.
- C. Use AWS DataSync to transfer the sequencing data to Amazon S3. Use S3 events to trigger an AWS Lambda function that starts an AWS Step Functions workflow. Store the Docker images in Amazon Elastic Container Registry (Amazon ECR) and trigger AWS Batch to run the container and process the sequencing data.
- D. Use an AWS Storage Gateway file gateway to transfer the sequencing data to Amazon S3. Use S3 events to trigger an AWS Batch job that runs on Amazon EC2 instances running the Docker containers to process the data.

Answer: C

NEW QUESTION 50

- (Exam Topic 2)

A company is in the process of implementing AWS Organizations to constrain its developers to use only Amazon EC2, Amazon S3, and Amazon DynamoDB. The developers' account resides in a dedicated organizational unit (OU). The solutions architect has implemented the following SCP on the developers' account:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowEC2",
      "Effect": "Allow",
      "Action": "ec2:*",
      "Resource": "*"
    },
    {
      "Sid": "AllowDynamoDB",
      "Effect": "Allow",
      "Action": "dynamodb:*",
      "Resource": "*"
    },
    {
      "Sid": "AllowS3",
      "Effect": "Allow",
      "Action": "s3:*",
      "Resource": "*"
    }
  ]
}
```

When this policy is deployed, IAM users in the developers' account are still able to use AWS services that are not listed in the policy. What should the solutions architect do to eliminate the developers' ability to use services outside the scope of this policy?

- A. Create an explicit deny statement for each AWS service that should be constrained.
- B. Remove the Full AWS Access SCP from the developer account's OU.
- C. Modify the Full AWS Access SCP to explicitly deny all services.
- D. Add an explicit deny statement using a wildcard to the end of the SCP.

Answer: B

NEW QUESTION 53

- (Exam Topic 2)

A company hosts a blog post application on AWS using Amazon API Gateway, Amazon DynamoDB, and AWS Lambda. The application currently does not use API keys to authorize requests. The API model is as follows:

GET /posts/{postId} to get post details
 GET /users/{userId}. to get user details
 GET /comments/{commentId}: to get comments details

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

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

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

Answer: C

NEW QUESTION 55

- (Exam Topic 2)

A company's solution architect is designing a disaster recovery (DR) solution for an application that runs on AWS. The application uses PostgreSQL 11.7 as its database. The company has an RPO of 30 seconds. The solution architect must design a DR solution with the primary database in the us-east-1 Region and the database in the us-west-2 Region.

What should the solution architect do to meet these requirements with minimum application change?

- A. Migrate the database to Amazon RDS for PostgreSQL in us-east-1. Set up a read replica in us-west-2. Set the managed RPO for the RDS database to 30 seconds.
- B. Migrate the database to Amazon RDS for PostgreSQL in us-east-1. Set up a standby replica in an Availability Zone in us-west-2. Set the managed RPO for the RDS database to 30 seconds.
- C. Migrate the database to an Amazon Aurora PostgreSQL global database with the primary Region as us-east-1 and the secondary Region as us-west-2. Set the managed RPO for the Aurora database to 30 seconds.
- D. Migrate the database to Amazon DynamoDB in us-east-1. Set up global tables with replica tables that are created in us-west-2.

Answer: A

NEW QUESTION 56

- (Exam Topic 2)

A company is using an Amazon CloudFront distribution to distribute both static and dynamic content from a web application running behind an Application Load Balancer. The web application requires user authorization and session tracking for dynamic content. The CloudFront distribution has a single cache behavior configured to forward the Authorization, Host, and Agent HTTP allow list headers and a session cookie to the origin. All other cache behavior settings are set to their default value.

A valid ACM certificate is applied to the CloudFront distribution with a matching CNAME in the distribution settings. The ACM certificate is also applied to the HTTPS listener for the Application Load Balancer. The CloudFront origin protocol policy is set to HTTPS only. Analysis of the cache statistics report shows that the miss rate for this distribution is very high.

What can the solution architect do to improve the cache hit rate for this distribution without causing the SSL/TLS handshake between CloudFront and the Application Load Balancer to fail?

- A. Create two cache behaviors for static and dynamic content. Remove the user-Agent and Host HTTP headers from the allow list headers section on both of the cache behaviors. Remove the session cookie from the allow list cookies section and the Authorization HTTP header from the allow list headers section for cache behavior configured for static content.
- B. Remove the user-Agent and Authorization HTTP headers from the allow list headers section of the cache behavior.
- C. Then update the cache behavior to use signed cookies for authorization.
- D. Remove the Host HTTP header from the allow list headers section and remove the session cookie from the allow list cookies section for the default cache behavior. Enable automatic object compression and use Lambda@Edge viewer request events for user authorization.
- E. Create two cache behaviors for static and dynamic content. Remove the User-Agent HTTP header from the allow list headers section on both of the cache behaviors. Remove the session cookie from the allow list cookies section and the Authorization HTTP header from the allow list headers section for cache behavior configured for static content.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/understanding-the-cache-key.html> Removing the host header will result in failed flow between CloudFront and ALB, because they have same certificate.

NEW QUESTION 60

- (Exam Topic 2)

A company wants to migrate its workloads from on-premises to AWS. The workloads run on Linux and Windows. The company has a large on-premises infrastructure that consists of physical machines and VMs that host numerous applications.

The company must capture details about the system configuration, system performance, running processes and network connections of its on-premises servers. The company also must divide the on-premises applications into groups for AWS migrations. The company needs recommendations for Amazon EC2 instance types so that the company can run its workloads on AWS in the most cost-effective manner.

Which combination of steps should a solution architect take to meet these requirements? (Select THREE.)

- A. Assess the existing applications by installing AWS Application Discovery Agent on the physical machines and VMs.
- B. Assess the existing applications by installing AWS Systems Manager Agent on the physical machines and VMs.
- C. Group servers into applications for migration by using AWS Systems Manager Application Manager.
- D. Group servers into applications for migration by using AWS Migration Hub.
- E. Generate recommended instance types and associated costs by using AWS Migration Hub.
- F. Import data about server sizes into AWS Trusted Advisor.
- G. Follow the recommendations for cost optimization.

Answer: BDF

NEW QUESTION 61

- (Exam Topic 2)

A company has more than 10,000 sensors that send data to an on-premises Apache Kafka server by using the Message Queuing Telemetry Transport (MQTT) protocol. The on-premises Kafka server transforms the data and then stores the results as objects in an Amazon S3 bucket.

Recently, the Kafka server crashed. The company lost sensor data while the server was being restored. A solution architect must create a new design on AWS that is highly available and scalable to prevent a similar occurrence.

Which solution will meet these requirements?

- A. Launch two Amazon EC2 instances to host the Kafka server in an active/standby configuration across two Availability Zones.
- B. Create a domain name in Amazon Route 53. Create a Route 53 failover policy. Route the sensors to send the data to the domain name.
- C. Migrate the on-premises Kafka server to Amazon Managed Streaming for Apache Kafka (Amazon MSK). Create a Network Load Balancer (NLB) that points to the Amazon MSK broker.
- D. Enable NLB health checks. Route the sensors to send the data to the NLB.
- E. Deploy AWS IoT Core, and connect it to an Amazon Kinesis Data Firehose delivery stream. Use an AWS Lambda function to handle data transformation. Route the sensors to send the data to AWS IoT Core.
- F. Deploy AWS IoT Core, and launch an Amazon EC2 instance to host the Kafka server. Configure AWS IoT Core to send the data to the EC2 instance. Route the sensors to send the data to AWS IoT Core.

Answer: A

NEW QUESTION 62

- (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 bucket
- 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. and 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 65

- (Exam Topic 2)

A company has an organization in AWS Organizations that has a large number of AWS accounts. One of the AWS accounts is designated as a transit account and has a transit gateway that is shared with all of the other AWS accounts. AWS Site-to-Site VPN connections are configured between all of the company's global offices and the transit account. The company has AWS Config enabled on all of its accounts.

The company's networking team needs to centrally manage a list of internal IP address ranges that belong to the global offices. Developers will reference this list to gain access to applications securely.

Which solution meets these requirements with the LEAST amount of operational overhead?

- A. Create a JSON file that is hosted in Amazon S3 and that lists all of the internal IP address ranges. Configure an Amazon Simple Notification Service (Amazon SNS) topic in each of the accounts that can be involved when the JSON file is updated.
- B. Subscribe an AWS Lambda function to the SNS topic to update all relevant security group rules with the updated IP address ranges.
- C. Create a new AWS Config managed rule that contains all of the internal IP address ranges. Use the rule to check the security groups in each of the accounts to ensure compliance with the list of IP address range.
- D. Configure the rule to automatically remediate any noncompliant security group that is detected.
- E. In the transit account, create a VPC prefix list with all of the internal IP address range.
- F. Use AWS Resource Access Manager to share the prefix list with all of the other accounts.
- G. Use the shared prefix list to configure security group rules in the other accounts.
- H. In the transit account, create a security group with all of the internal IP address range.
- I. Configure the security groups in the other accounts to reference the transit account's security group by using a nested security group reference of `*-<transit-account-id>/.sg-1a2b3c4d`.

Answer: C

NEW QUESTION 68

- (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 group.
- 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 group.
- 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 71

- (Exam Topic 2)

A large company recently experienced an unexpected increase in Amazon RDS and Amazon DynamoDB costs. The company needs to increase visibility into details of AWS Billing and Cost Management. There are various accounts associated with AWS Organizations, including many development and production accounts. There is no consistent tagging strategy across the organization, but there are guidelines in place that require all infrastructure to be deployed using AWS CloudFormation with consistent tagging. Management requires cost center numbers and project ID numbers for all existing and future DynamoDB tables and RDS instances.

Which strategy should the solutions architect provide to meet these requirements?

- A. Use Tag Editor to tag existing resources Create cost allocation tags to define the cost center and project ID and allow 24 hours for tags to propagate to existing resources
- B. Use an AWS Config rule to alert the finance team of untagged resources Create a centralized AWS Lambda based solution to tag untagged RDS databases and DynamoDB resources every hour using a cross-account role.
- C. Use Tag Editor to tag existing resources Create cost allocation tags to define the cost center and project ID Use SCPs to restrict resource creation that do not have the cost center and project ID on the resource.
- D. Create cost allocation tags to define the cost center and project ID and allow 24 hours for tags to propagate to existing resources Update existing federated roles to restrict privileges to provision resources that do not include the cost center and project ID on the resource

Answer: B

NEW QUESTION 76

- (Exam Topic 2)

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

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

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

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

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

Answer: BDE

NEW QUESTION 81

- (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 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 S3Cross-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 S3Same-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: B

NEW QUESTION 86

- (Exam Topic 2)

A company is configuring connectivity to a multi-account AWS environment to support application workloads that serve users in a single geographic region. The workloads depend on a highly available, on-premises legacy system deployed across two locations It is critical for the AWS workloads to maintain connectivity to the legacy system, and a minimum of 5 Gbps of bandwidth is required All application workloads within AWS must have connectivity with one another.

Which solution will meet these requirements?

- A. Configure multiple AWS Direct Connect (DX) 10 Gbps dedicated connections from a DX partner for each on-premises location Create private virtual interfaces on each connection for each AWS account VPC Associate the private virtual interface with a virtual private gateway attached to each VPC
- B. Configure multiple AWS Direct Connect (DX) 10 Gbps dedicated connections from two DX partners for each on-premises location Create and attach a virtual private gateway for each AWS account VPC
- C. Create a DX gateway in a central network account and associate it with the virtual private gateways Create a public virtual interface on each DX connection and associate the interface with the DX gateway.
- D. Configure multiple AWS Direct Connect (DX) 10 Gbps dedicated connections from two DX partners for each on-premises location Create a transit gateway and a DX gateway in a central network account
- E. Create a transit virtual interface for each DX connection and associate them with the DX gateway
- F. Create a gateway association between the DX gateway and the transit gateway
- G. Configure multiple AWS Direct Connect (DX) 10 Gbps dedicated connections from a DX partner for each on-premises location Create and attach a virtual private gateway for each AWS account VPC
- H. Create a transit gateway in a central network account and associate it with the virtual private gateways Create a transit virtual interface on each DX connection and attach the interface to the transit gateway.

Answer: B

NEW QUESTION 90

- (Exam Topic 2)

A company's CI SO has asked a solutions architect to re-engineer the company's current CI/CD practices to make sure patch deployments to its application can happen as quickly as possible with minimal downtime if vulnerabilities are discovered. The company must also be able to quickly roll back a change in case of errors.

The web application is deployed in a fleet of Amazon EC2 instances behind an Application Load Balancer. The company is currently using GitHub to host the application source code, and has configured an AWS CodeBuild project to build the application. The company also intends to use AWS CodePipeline to trigger builds from GitHub commits using the existing CodeBuild project.

What CI/CD configuration meets all of the requirements?

- A. Configure CodePipeline with a deploy stage using AWS CodeDeploy configured for in-place deployment. Monitor the newly deployed code, and, if there are any issues, push another code update.
- B. Configure CodePipeline with a deploy stage using AWS CodeDeploy configured for blue/green deployments. Monitor the newly deployed code and, if there are any issues, trigger a manual rollback using CodeDeploy.
- C. Configure CodePipeline with a deploy stage using AWS CloudFormation to create a pipeline for test and production stacks. Monitor the newly deployed code, and, if there are any issues, push another code update.
- D. Configure the CodePipeline with a deploy stage using AWS OpsWorks and in-place deployments. Monitor the newly deployed code and, if there are any issues, push another code update.
- E. If there are any issues, push another code update.

Answer: B

NEW QUESTION 92

- (Exam Topic 2)

A company is using multiple AWS accounts. The DNS records are stored in a private hosted zone for Amazon Route 53 in Account A. The company's applications and databases are running in Account B.

A solutions architect will deploy a two-tier application in a new VPC. To simplify the configuration, the db.example.com CNAME record set for the Amazon RDS endpoint was created in a private hosted zone for Amazon Route 53.

During deployment, the application failed to start. Troubleshooting revealed that db.example.com is not resolvable on the Amazon EC2 instance. The solutions architect confirmed that the record set was created correctly in Route 53.

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

- A. Deploy the database on a separate EC2 instance in the new VPC. Create a record set for the instance's private IP in the private hosted zone.
- B. Use SSH to connect to the application tier EC2 instance. Add an RDS endpoint IP address to the /etc/resolv.conf file.
- C. Create an authorization to associate the private hosted zone in Account A with the new VPC in Account B.
- D. Create a private hosted zone for the example.com domain in Account B. Configure Route 53 replication between AWS accounts.
- E. Associate a new VPC in Account B with a hosted zone in Account A.
- F. Delete the association authorization in Account A.

Answer: CE

NEW QUESTION 94

- (Exam Topic 2)

A company's solutions architect is reviewing a web application that runs on AWS. The application references static assets in an Amazon S3 bucket in the us-east-1 Region. The company needs resiliency across multiple AWS Regions. The company already has created an S3 bucket in a second Region.

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

- A. Configure the application to write each object to both S3 buckets.
- B. Set up an Amazon Route 53 public hosted zone with a record set by using a weighted routing policy for each S3 bucket.
- C. Configure the application to reference the objects by using the Route 53 DNS name.
- D. Create an AWS Lambda function to copy objects from the S3 bucket in us-east-1 to the S3 bucket in the second Region.
- E. Invoke the Lambda function each time an object is written to the S3 bucket in us-east-1. Set up an Amazon CloudFront distribution with an origin group that contains the two S3 buckets as origins.
- F. Configure replication on the S3 bucket in us-east-1 to replicate objects to the S3 bucket in the second Region. Set up an Amazon CloudFront distribution with an origin group that contains the two S3 buckets as origins.
- G. Configure replication on the S3 bucket in us-east-1 to replicate objects to the S3 bucket in the second Region.
- H. If failover is required, update the application code to load S3 objects from the S3 bucket in the second Region.

Answer: D

NEW QUESTION 99

- (Exam Topic 2)

A company runs an IoT platform on AWS. IoT sensors in various locations send data to the company's Node.js API servers on Amazon EC2 instances running behind an Application Load Balancer. The data is stored in an Amazon RDS MySQL DB instance that uses a 4 TB General Purpose SSD volume.

The number of sensors the company has deployed in the field has increased over time and is expected to grow significantly. The API servers are consistently overloaded, and RDS metrics show high write latency.

Which of the following steps together will resolve the issues permanently and enable growth as new sensors are provisioned, while keeping this platform cost-efficient? (Select TWO.)

- A. Resize the MySQL General Purpose SSD storage to 6 TB to improve the volume's IOPS.
- B. Re-architect the database tier to use Amazon Aurora instead of an RDS MySQL DB instance and add read replicas.
- C. Leverage Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data.
- D. Use AWS X-Ray to analyze and debug application issues and add more API servers to match the load.
- E. Re-architect the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance.

Answer: CE

NEW QUESTION 101

- (Exam Topic 2)

A company has an organization in AWS Organizations. The organization consists of a large number of AWS accounts that belong to separate business units. The

company requires all Amazon EC2 instances to be provisioned with custom, hardened AMIs. The company wants a solution that provides each AWS account access to the AMIs

Which solution will meet these requirements with the MOST operational efficiency?

- A. Create the AMIs with EC2 Image Builder Create an AWS CodePipeline pipeline to share the AMIs across all AWS accounts.
- B. Deploy Jenkins on an EC2 instance Create jobs to create and share the AMIs across all AWS accounts.
- C. Create and share the AMIs with EC2 Image Builder Use AWS Service Catalog to configure a product that provides access to the AMIs across all AWS accounts.
- D. Create the AMIs with EC2 Image Builder Create an AWS Lambda function to share the AMIs across all AWS accounts.

Answer: C

NEW QUESTION 103

- (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 104

- (Exam Topic 2)

A financial services company in North America plans to release a new online web application to its customers on AWS. The company will launch the application in the us-east-1 Region on Amazon EC2 instances. The application must be highly available and must dynamically scale to meet user traffic. The company also wants to implement a disaster recovery environment for the application in the us-west-1 Region by using active-passive failover.

Which solution will meet these requirements?

- A. Create a VPC in us-east-1 and a VPC in us-west-1. Configure VPC peering in the us-east-1 VPC.
- B. Create an Application Load Balancer (ALB) that extends across multiple Availability Zones in both VPCs. Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in both VPCs. Place the Auto Scaling group behind the ALB.
- C. Create a VPC in us-east-1 and a VPC in us-west-1. In the us-east-1 VPC.
- D. Create an Application Load Balancer (ALB) that extends across multiple Availability Zones in that VPC.
- E. Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in the us-east-1 VPC. Place the Auto Scaling group behind the ALB. Set up the same configuration in the us-west-1 VPC.
- F. Create an Amazon Route 53 hosted zone. Create separate records for each ALB. Enable health checks to ensure high availability between Regions.
- G. Create a VPC in us-east-1 and a VPC in us-west-1. In the us-east-1 VPC.
- H. Create an Application Load Balancer (ALB) that extends across multiple Availability Zones in that VPC. Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in the us-east-1 VPC. Place the Auto Scaling group behind the ALB. Set up the same configuration in the us-west-1 VPC. Create an Amazon Route 53 hosted zone.
- I. Create separate records for each ALB. Enable health checks and configure a failover routing policy for each record.
- J. Create a VPC in us-east-1 and a VPC in us-west-1. Configure VPC peering in the us-east-1 VPC.
- K. Create an Application Load Balancer (ALB) that extends across multiple Availability Zones. Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in both VPCs. Place the Auto Scaling group behind the ALB. Create an Amazon Route 53 host. Create a record for the ALB.

Answer: C

NEW QUESTION 106

- (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 theirbased on the ta
F. day, and time

Answer: C

NEW QUESTION 109

- (Exam Topic 2)

A company is running an application in the AWS Cloud. The application uses AWS Lambda functions and Amazon Elastic Container Service (Amazon ECS) containers that run with AWS Fargate technology as its primary compute. The load on the application is irregular. The application experiences long periods of no usage, followed by sudden and significant increases and decreases in traffic. The application is write-heavy and stores data in an Amazon Aurora MySQL database. The database runs on an Amazon RDS memory optimized D8 instance that is not able to handle the load.

What is the MOST cost-effective way for the company to handle the sudden and significant changes in traffic?

- A. Add additional read replicas to the databas
- B. Purchase Instance Savings Plans and RDS Reserved Instances.
- C. Migrate the database to an Aurora multi-master DB cluste
- D. Purchase Instance Savings Plans.
- E. Migrate the database to an Aurora global database Purchase Compute Savings Plans and RDS Reserved Instances
- F. Migrate the database to Aurora Serverless v1. Purchase Compute Savings Plans

Answer: D

NEW QUESTION 111

- (Exam Topic 2)

An ecommerce company runs its infrastructure on AWS. The company exposes its APIs to its web and mobile clients through an Application Load Balancer (ALB) in front of an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The EKS cluster runs thousands of pods that provide the APIs.

After extending delivery to a new continent, the company adds an Amazon CloudFront distribution and sets the ALB as the origin. The company also adds AWS WAF to its architecture.

After implementation of the new architecture, API calls are significantly. However, there is a sudden increase in HTTP status code 504 (Gateway Timeout) errors and HTTP status code 502 (Bad Gateway) errors. This increase in errors seems to be for a specific domain. Which factors could be a cause of these errors? (Select TWO.)

- A. AWS WAF is blocking suspicious requests.
- B. The origin is not properly configured in CloudFront.
- C. There is an SSL/TLS handshake issue between CloudFront and the origin.
- D. EKS Kubernetes pods are being cycled.
- E. Some pods are taking more than 30 seconds to answer API calls.

Answer: AE

NEW QUESTION 112

- (Exam Topic 2)

A software company has deployed an application that consumes a REST API by using Amazon API Gateway. AWS Lambda functions, and an Amazon DynamoDB table. The application is showing an increase in the number of errors during PUT requests. Most of the PUT calls come from a small number of clients that are authenticated with specific API keys.

A solutions architect has identified that a large number of the PUT requests originate from one client. The API is noncritical, and clients can tolerate retries of unsuccessful calls. However, the errors are displayed to customers and are causing damage to the API's reputation.

What should the solutions architect recommend to improve the customer experience?

- A. Implement retry logic with exponential backoff and irregular variation in the client applicatio
- B. Ensure that the errors are caught and handled with descriptive error messages.
- C. Implement API throttling through a usage plan at the API Gateway leve
- D. Ensure that the client application handles code 429 replies without error.
- E. Turn on API caching to enhance responsiveness for the production stag
- F. Run 10-minute load tests. Verify that the cache capacity is appropriate for the workload.
- G. Implement reserved concurrency at the Lambda function level to provide the resources that are needed during sudden increases in traffic.

Answer: A

NEW QUESTION 115

- (Exam Topic 2)

A company is running an application in the AWS Cloud. The application consists of microservices that run on a fleet of Amazon EC2 instances in multiple Availability Zones behind an Application Load Balancer. The company recently added a new REST API that was implemented in Amazon API Gateway. Some of the older microservices that run on EC2 instances need to call this new API

The company does not want the API to be accessible from the public internet and does not want proprietary data to traverse the public internet

What should a solutions architect do to meet these requirements?

- A. Create an AWS Site-to-Site VPN connection between the VPC and the API Gateway Use API Gateway to generate a unique API key for each microservic
- B. Configure the API methods to require the key.
- C. Create an interface VPC endpoint for API Gateway, and set an endpoint policy to only allow access to the specific API Add a resource policy to API Gateway to only allow access from the VPC endpoint Change the API Gateway endpoint type to private.
- D. Modify the API Gateway to use IAM authentication Update the IAM policy for the IAM role that is assigned to the EC2 instances to allow access to the API Gateway Move the API Gateway into a new VPC Deploy a transit gateway and connect the VPCs.
- E. Create an accelerator in AWS Global Accelerator and connect the accelerator to the API Gateway. Update the route table for all VPC subnets with a route to the created Global Accelerator endpoint IP address
- F. Add an API key for each service to use for authentication.

Answer: B

NEW QUESTION 118

- (Exam Topic 2)

A company wants to deploy an API to AWS. The company plans to run the API on AWS Fargate behind a load balancer. The API requires the use of header-based routing and must be accessible from on-premises networks through an AWS Direct Connect connection and a private VIF.

The company needs to add the client IP addresses that connect to the API to an allow list in AWS. The company also needs to add the IP addresses of the API to the allow list. The company's security team will allow /27 CIDR ranges to be added to the allow list. The solution must minimize complexity and operational overhead.

Which solution will meet these requirements?

- A. Create a new Network Load Balancer (NLB) in the same subnets as the Fargate task deployments. Create a security group that includes only the client IP addresses that need access to the AP
- B. Attach the new security group to the Fargate task
- C. Provide the security team with the NLB's IP addresses for the allow list.
- D. Create two new /27 subnet
- E. Create a new Application Load Balancer (ALB) that extends across the new subnet
- F. Create a security group that includes only the client IP addresses that need access to the AP
- G. Attach the security group to the AL
- H. Provide the security team with the new subnet IP ranges for the allow list.
- I. Create two new '27 subnet
- J. Create a new Network Load Balancer (NLB) that extends across the new subnet
- K. Create a new Application Load Balancer (ALB) within the new subnet
- L. Create a security group that includes only the client IP addresses that need access to the AP
- M. Attach the security group to the AL
- N. Add the ALB's IP addresses as targets behind the NL
- O. Provide the security team with the NLB's IP addresses for the allow list.
- P. Create a new Application Load Balancer (ALB) in the same subnets as the Fargate task deployments. Create a security group that includes only the client IP addresses that need access to the AP
- Q. Attach the security group to the AL
- R. Provide the security team with the ALB's IP addresses for the allow list.

Answer: A

NEW QUESTION 119

- (Exam Topic 2)

An AWS partner company is building a service in AWS Organizations using its organization named org. This service requires the partner company to have access to AWS resources in a customer account, which is in a separate organization named org2. The company must establish least privilege security access using an API or command line tool to the customer account.

What is the MOST secure way to allow org1 to access resources in org2?

- A. The customer should provide the partner company with their AWS account access keys to log in and perform the required tasks.
- B. The customer should create an IAM user and assign the required permissions to the IAM user. The customer should then provide the credentials to the partner company to log in and perform the required tasks.
- C. The customer should create an IAM role and assign the required permissions to the IAM role.
- D. The partner company should then use the IAM role's Amazon Resource Name (ARN) when requesting access to perform the required tasks.
- E. The customer should create an IAM role and assign the required permissions to the IAM role.
- F. The partner company should then use the IAM role's Amazon Resource Name (ARN). Including the external ID in the IAM role's trust policy, when requesting access to perform the required tasks.

Answer: D

NEW QUESTION 124

- (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 126

- (Exam Topic 2)

A company wants to improve cost awareness for its Amazon EMR platform. The company has allocated budgets for each team's Amazon EMR usage. When a budgetary threshold is reached, a notification should be sent by email to the budget office's distribution list. Teams should be able to view their EMR cluster expenses to date. A solutions architect needs to create a solution that ensures this policy is proactively and centrally enforced in a multi-account environment. Which combination of steps should the solutions architect take to meet these requirements? (Select TWO.)

- A. Update the AWS CloudFormation template to include the AWS Budgets Budget resource with the NotificationsWithSubscribers property.
- B. Implement Amazon CloudWatch dashboards for Amazon EMR usage.
- C. Create an EMR bootstrap action that runs at startup that calls the Cost Explorer API to set the budget on the cluster with the GetCostForecast and NotificationsWithSubscribers actions.
- D. Create an AWS Service Catalog portfolio for each team.

- E. Add each team's Amazon EMR cluster as an AWS CloudFormation template to their Service Catalog portfolio as a Product
- F. Create an Amazon CloudWatch metric for billing Create a custom alert when costs exceed the budgetary threshold.

Answer: BE

NEW QUESTION 127

- (Exam Topic 2)

A company is using an Amazon EMR cluster to run its big data jobs The cluster's jobs are invoked by AWS

Step Functions Express Workflows that consume various Amazon Simple Queue Service (Amazon SQS) queues The workload of this solution is variable and unpredictable Amazon CloudWatch metrics show that the cluster's peak utilization is only 25% at times and that the cluster sits idle the rest of the time

A solutions architect must optimize the costs of the cluster without negatively impacting the time it takes to run the various jobs

What is the MOST cost-effective solution that meets these requirements?

- A. Modify the EMR cluster by turning on automatic scaling of the core nodes and task nodes with a custom policy that is based on cluster utilization Purchase Reserved Instance capacity to cover the master node.
- B. Modify the EMR cluster to use an instance fleet of Dedicated On-Demand Instances for the master node and core nodes, and to use Spot Instances for the task node
- C. Define target capacity for each node type to cover the load.
- D. Purchase Reserved Instances for the master node and core nodes Terminate all existing task nodes in the EMR cluster
- E. Modify the EMR cluster to use capacity-optimized Spot Instances and a diversified task fleet
- F. Define target capacity for each node type with a mix of On-Demand Instances and Spot Instances.

Answer: B

NEW QUESTION 130

- (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 all 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 133

- (Exam Topic 2)

An auction website enables users to bid on collectible items The auction rules require that each bid is processed only once and in the order it was received The current implementation is based on a fleet of Amazon EC2 web servers that write bid records into Amazon Kinesis Data Streams A single large instance has a cron job that runs the bid processor, which reads incoming bids from Kinesis Data Streams and processes each bid The auction site is growing in popularity, but users are complaining that some bids are not registering

Troubleshooting indicates that the bid processor is too slow during peak demand hours sometimes crashes while processing and occasionally loses track of which record is being processed

What changes should make the bid processing more reliable?

- A. Refactor the web application to use the Amazon Kinesis Producer Library (KPL) when posting bids to Kinesis Data Streams Refactor the bid processor to flag each record in Kinesis Data Streams as being unread processing and processed At the start of each bid processing run; scan Kinesis Data Streams for unprocessed records
- B. Refactor the web application to post each incoming bid to an Amazon SNS topic in place of Kinesis Data Streams Configure the SNS topic to trigger an AWS Lambda function that
- C. processes each bid as soon as a user submits it
- D. Refactor the web application to post each incoming bid to an Amazon SQS FIFO queue in place of Kinesis Data Streams Refactor the bid processor to continuously consume the SQS queue Place the bid processing EC2 instance in an Auto Scaling group with a minimum and a maximum size of 1
- E. Switch the EC2 instance type from t2 large to a larger general compute instance type Put the bid processor EC2 instances in an Auto Scaling group that scales out the number of EC2 instances running the bid processor based on the incomingRecords metric in Kinesis Data Streams

Answer: C

Explanation:

<https://aws.amazon.com/sqs/faqs/#:~:text=A%20single%20Amazon%20SQS%20message,20%2C000%20for%2>

NEW QUESTION 135

- (Exam Topic 2)

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

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

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

```

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

```

D. In Account A, set the S3 bucket policy to the following: Text, letter Description automatically generated

```

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

```

E. In Account B, set the permissions of User_DataProcessor to the following: Text Description automatically generated

```

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

```

F. In Account B, set the permissions of User_DataProcessor to the following: Text, letter Description automatically generated

```

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

```

Answer: AD

NEW QUESTION 137

- (Exam Topic 2)

During an audit, a security team discovered that a development team was putting IAM user secret access keys in their code and then committing it to an AWS CodeCommit repository . The security team wants to automatically find and remediate instances of this security vulnerability Which solution will ensure that the credentials are appropriately secured automatically?

- A. Run a script nightly using AWS Systems Manager Run Command to search for credentials on the development instances If found use AWS Secrets Manager to rotate the credentials.
- B. Use a scheduled AWS Lambda function to download and scan the application code from CodeCommit If credentials are found, generate new credentials and store them in AWS KMS
- C. Configure Amazon Macie to scan for credentials in CodeCommit repositories If credentials are found, trigger an AWS Lambda function to disable the credentials and notify the user
- D. Configure a CodeCommit trigger to invoke an AWS Lambda function to scan new code submissions for credentials If credentials are found, disable them in AWS IAM and notify the user.

Answer: A

NEW QUESTION 141

- (Exam Topic 2)

A solutions architect is designing a solution to connect a company's on-premises network with all the company's current and future VPCs on AWS The company is running VPCs in five different AWS Regions and has at least 15 VPCs in each Region.

The company's AWS usage is constantly increasing and will continue to grow Additionally, all the VPCs throughout all five Regions must be able to communicate with each other

The solution must maximize scalability and ease of management Which solution meets these requirements'?

- A. Set up a transit gateway in each Region Establish a redundant AWS Site-to-Site VPN connection between the on-premises firewalls and the transit gateway in the Region that is closest to the on-premises network Peer all the transit gateways with each other Connect all the VPCs to the transit gateway in their Region
- B. Create an AWS CloudFormation template for a redundant AWS Site-to-Site VPN tunnel to the on-premises network Deploy the CloudFormation template for each VPC Set up VPC peering between all the VPCs for VPC-to-VPC communication
- C. Set up a transit gateway in each Region Establish a redundant AWS Site-to-Site VPN connection between the on-premises firewalls and each transit gateway Route traffic between the different Regions through the company's on-premises firewalls Connect all the VPCs to the transit gateway in their Region
- D. Create an AWS CloudFormation template for a redundant AWS Site-to-Site VPN tunnel to the on-premises network Deploy the CloudFormation template for each VPC Route traffic between the different Regions through the company's on-premises firewalls

Answer: A

NEW QUESTION 145

- (Exam Topic 2)

A solutions architect is importing a VM from an on-premises environment by using the Amazon EC2 VM Import feature of AWS Import/Export The solutions architect has created an AMI and has provisioned an Amazon EC2 instance that is based on that AMI The EC2 instance runs inside a public subnet in a VPC and has a public IP address assigned

The EC2 instance does not appear as a managed instance in the AWS Systems Manager console

Which combination of steps should the solutions architect take to troubleshoot this issue"? (Select TWO)

- A. Verify that Systems Manager Agent is installed on the instance and is running
- B. Verify that the instance is assigned an appropriate IAM role for Systems Manager
- C. Verify the existence of a VPC endpoint on the VPC
- D. Verify that the AWS Application Discovery Agent is configured
- E. Verify the correct configuration of service-linked roles for Systems Manager

Answer: ABD

NEW QUESTION 149

- (Exam Topic 2)

A company is collecting a large amount of data from a fleet of IoT devices. Data is stored as Optimized Row Columnar (ORC) files in the Hadoop Distributed File System (HDFS) on a persistent Amazon EMR cluster. The company's data analytics team queries the data by using SQL in Apache Presto deployed on the same EMR cluster. Queries scan large amounts of data always run for less than 15 minutes, and run only between 5 PM and 10 PM.

The company is concerned about the high cost associated with the current solution. A solutions architect must propose the most cost-effective solution that will allow SQL data queries.

Which solution will meet these requirements?

- A. Store data in Amazon S3. Use Amazon Redshift Spectrum to query data.
- B. Store data in Amazon S3. Use the AWS Glue Data Catalog and Amazon Athena to query data.
- C. Store data in EMR File System (EMRFS). Use Presto on Amazon EMR to query data.
- D. Store data in Amazon Redshift. Use Amazon Redshift to query data.

Answer: D

NEW QUESTION 153

- (Exam Topic 2)

A company is using an existing orchestration tool to manage thousands of Amazon EC2 instances. A recent penetration test found a vulnerability in the company's software stack. This vulnerability has prompted the company to perform a full evaluation of its current production environment. The analysts determined that the following vulnerabilities exist within the environment:

- Operating systems with outdated libraries and known vulnerabilities are being used in production
- Relational databases hosted and managed by the company are running unsupported versions with known vulnerabilities
- Data stored in databases is not encrypted.

The solutions architect intends to use AWS Config to continuously audit and assess the compliance of the company's AWS resource configurations with the company's policies and guidelines. What additional steps will enable the company to secure its environments and track resources while adhering to best practices?

- A. Use AWS Application Discovery Service to evaluate all running EC2 instances. Use the AWS CLI to modify each instance, and use EC2 user data to install the AWS Systems Manager Agent during boot. Schedule patching to run as a Systems Manager Maintenance Window.
- B. Migrate all relational databases to Amazon RDS and enable AWS KMS encryption.
- C. Create an AWS CloudFormation template for the EC2 instances. Use EC2 user data in the CloudFormation template to install the AWS Systems Manager Agent, and enable AWS KMS encryption on all Amazon EBS volumes.
- D. Have CloudFormation replace all running instances.
- E. Use Systems Manager Patch Manager to establish a patch baseline and deploy a Systems Manager Maintenance Windows task to run AWS-RunPatchBaseline using the patch baseline.
- F. Install the AWS Systems Manager Agent on all existing instances using the company's current orchestration tool. Use the Systems Manager Run Command to run a list of commands to upgrade software on each instance using operating system-specific tools.
- G. Enable AWS KMS encryption on all Amazon EBS volumes.
- H. Install the AWS Systems Manager Agent on all existing instances using the company's current orchestration tool.
- I. Migrate all relational databases to Amazon RDS and enable AWS KMS encryption. Use Systems Manager Patch Manager to establish a patch baseline and deploy a Systems Manager Maintenance Windows task to run AWS-RunPatchBaseline using the patch baseline.

Answer: D

NEW QUESTION 156

- (Exam Topic 2)

A new application is running on Amazon Elastic Container Service (Amazon ECS) with AWS Fargate. The application uses an Amazon Aurora MySQL database. The application and the database run in the same subnets of a VPC with distinct security groups that are configured.

The password (or the database is stored in AWS Secrets Manager and is passed to the application through the `D8_PASSWORD` environment variable. The hostname of the database is passed to the application through the `DB_HOST` environment variable. The application is failing to access the database.

Which combination of actions should a solutions architect take to resolve this error? (Select THREE.)

- A. Ensure that the container has the environment variable with name `"DB_PASSWORD"` specified with a `"ValueFrom"` and the ARN of the secret.
- B. Ensure that the container has the environment variable with name `"D8_PASSWORD"` specified with a `"ValueFrom"` and the secret name of the secret.
- C. Ensure that the Fargate service security group allows inbound network traffic from the Aurora MySQL database on the MySQL TCP port 3306.
- D. Ensure that the Aurora MySQL database security group allows inbound network traffic from the Fargate service on the MySQL TCP port 3306.
- E. Ensure that the container has the environment variable with name `"D8_HOST"` specified with the hostname of a DB instance endpoint.
- F. Ensure that the container has the environment variable with name `"DB_HOST"` specified with the hostname of the DB instance endpoint.

Answer: ADE

NEW QUESTION 161

- (Exam Topic 2)

A large company has a business-critical application that runs in a single AWS Region. The application consists of multiple Amazon EC2 instances and an Amazon RDS Multi-AZ DB instance. The EC2 instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones.

A solutions architect is implementing a disaster recovery (DR) plan for the application. The solutions architect has created a pilot light application deployment in a new Region, which is referred to as the DR Region. The DR environment has an Auto Scaling group with a single EC2 instance and a read replica of the RDS DB instance.

The solutions architect must automate a failover from the primary application environment to the pilot light environment in the DR Region.

Which solution meets these requirements with the MOST operational efficiency?

- A. Publish an application availability metric to Amazon CloudWatch in the DR Region from the application environment in the primary Region. Create a CloudWatch alarm in the DR Region that is invoked when the application availability metric stops being delivered. Configure the CloudWatch alarm to send a notification to an Amazon Simple Notification Service (Amazon SNS) topic in the DR Region. Add an email subscription to the SNS topic that sends messages to the application owner upon notification, instruct a systems operator to sign in to the AWS Management Console and initiate failover operations for the application.
- B. Create a cron task that runs every 5 minutes by using one of the application's EC2 instances in the primary Region. Configure the cron task to check whether the application is available. Upon failure, the cron task notifies a systems operator and attempts to restart the application services.

- C. Create a cron task that runs every 5 minutes by using one of the application's EC2 instances in the primary Region Configure the cron task to check whether the application is available Upon failure, the cron task modifies the DR environment by promoting the read replica and by adding EC2 instances to the Auto Scaling group
- D. Publish an application availability metric to Amazon CloudWatch in the DR Region from the application environment in the primary Region Create a CloudWatch alarm in the DR Region that is invoked when the application availability metric stops being delivered Configure the CloudWatch alarm to send a notification to an Amazon Simple Notification Service (Amazon SNS) topic in the DR Region Use an AWS Lambda function that is invoked by Amazon SNS in the DR Region to promote the read replica and to add EC2 instances to the Auto Scaling group

Answer: D

NEW QUESTION 163

- (Exam Topic 2)

A company is running an application in the AWS Cloud. The company has several third-party services that integrate with the application through a RESTful API. The API is a serverless implementation with an Amazon API Gateway regional API endpoint that integrates with several different AWS Lambda functions. The application's data is nonrelational and is stored in an Amazon DynamoDB table. The application and the API are running in the eu-west-1 Region. The company needs the API to also be available in the us-east-1 Region. All data must be available in both Regions. A solutions architect already has deployed all the Lambda functions in us-east-1

Which additional steps should the solutions architect take to meet these requirements? (Select TWO.)

- A. Deploy a second API Gateway regional API endpoint in us-east-1. Create Lambda integration with the functions in us-east-1.
- B. Enable DynamoDB Streams on the table in eu-west-1. Replicate all changes to a DynamoDB table in us-east-1
- C. Modify the DynamoDB table to be a global table in eu-west-1 and in us-east-1.
- D. Change the API Gateway API endpoint in eu-west-1 to an edge-optimized endpoint
- E. Create Lambda integration with the functions in both Regions.
- F. Create a DynamoDB read replica in us-east-1.

Answer: AC

NEW QUESTION 166

- (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 171

- (Exam Topic 2)

A company has automated the nightly retraining of its machine learning models by using AWS Step Functions. The workflow consists of multiple steps that use AWS Lambda. Each step can fail for various reasons, and any failure causes a failure of the overall workflow.

A review reveals that the retraining has failed multiple nights in a row without the company noticing the failure. A solutions architect needs to improve the workflow so that notifications are sent for all types of failures in the retraining process.

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

- A. Create an Amazon Simple Notification Service (Amazon SNS) topic with a subscription of type "Email" that targets the team's mailing list.
- B. Create a task named "Email" that forwards the input arguments to the SNS topic
- C. Add a Catch field to all Tasks
- D. Map
- E. and Parallel states that have a statement of "ErrorEquals": ["states.all"] and "Next": "Email".
- F. Add a new email address to Amazon Simple Email Service (Amazon SES). Verify the email address.
- G. Create a task named "Email" that forwards the input arguments to the SES email address
- H. Add a Catch field to all Task, Map, and Parallel states that have a statement of "ErrorEquals": ["states.Bun time"] and "Next": "Email".

Answer: BCD

NEW QUESTION 175

- (Exam Topic 2)

A company is currently using AWS CodeCommit for its source control and AWS CodePipeline for continuous integration The pipeline has a build stage for building the artifacts, which is then staged in an Amazon S3 bucket.

The company has identified various improvement opportunities in the existing process and a solutions architect has been given the following requirements

- Create a new pipeline to support feature development
- Support feature development without impacting production applications
- Incorporate continuous testing with unit tests
- Isolate development and production artifacts

- Support the capability to merge tested code into production code How should the solutions architect achieve these requirements?

- A. Trigger a separate pipeline from CodeCommit feature branches Use AWS CodeBuild for running unit tests Use CodeBuild to stage the artifacts within an S3 bucket in a separate testing account
- B. Trigger a separate pipeline from CodeCommit feature branches Use AWS Lambda for running unit tests Use AWS CodeDeploy to stage the artifacts within an S3 bucket in a separate testing account
- C. Trigger a separate pipeline from CodeCommit tags Use Jenkins for running unit tests Create a stage in the pipeline with S3 as the target for staging the artifacts within an S3 bucket in a separate testing account.
- D. Create a separate CodeCommit repository for feature development and use it to trigger the pipeline Use AWS Lambda for running unit tests Use AWS CodeBuild to stage the artifacts within different S3 buckets in the same production account

Answer: A

NEW QUESTION 179

- (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 180

- (Exam Topic 2)

A medical company is running an application in the AWS Cloud. The application simulates the effect of medical drugs in development. The application consists of two parts: configuration and simulation. The configuration part runs in AWS Fargate containers in an Amazon Elastic Container Service (Amazon ECS) cluster. The simulation part runs on large, compute-optimized Amazon EC2 instances. Simulations can restart if they are interrupted. The configuration part runs 24 hours a day with a steady load. The simulation part runs only for a few hours each night with a variable load. The company stores simulation results in Amazon S3, and researchers use the results for 30 days. The company must store simulations for 10 years and must be able to retrieve the simulations within 5 hours.

Which solution meets these requirements MOST cost-effectively?

- A. Purchase an EC2 Instance Savings Plan to cover the usage for the configuration part. Run the simulation part by using EC2 Spot Instances. Create an S3 Lifecycle policy to transition objects that are older than 30 days to S3 Intelligent-Tiering.
- B. Purchase an EC2 Instance Savings Plan to cover the usage for the configuration part and the simulation part. Create an S3 Lifecycle policy to transition objects that are older than 30 days to S3 Glacier.
- C. Purchase Compute Savings Plans to cover the usage for the configuration part. Run the simulation part by using EC2 Spot instances. Create an S3 Lifecycle policy to transition objects that are older than 30 days to S3 Glacier.
- D. Purchase Compute Savings Plans to cover the usage for the configuration part. Purchase EC2 Reserved Instances for the simulation part. Create an S3 Lifecycle policy to transition objects that are older than 30 days to S3 Glacier Deep Archive.

Answer: C

NEW QUESTION 181

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

SAP-C02 Practice Exam Features:

- * SAP-C02 Questions and Answers Updated Frequently
- * SAP-C02 Practice Questions Verified by Expert Senior Certified Staff
- * SAP-C02 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * SAP-C02 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The SAP-C02 Practice Test Here](#)