

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional

<https://www.2passeasy.com/dumps/AWS-Certified-Solutions-Architect-Professional/>



NEW QUESTION 1

- (Exam Topic 2)

An external audit of a company's serverless application reveals IAM policies that grant too many permissions. These policies are attached to the company's AWS Lambda execution roles. Hundreds of the company's Lambda functions have broad access permissions, such as full access to Amazon S3 buckets and Amazon DynamoDB tables. The company wants each function to have only the minimum permissions that the function needs to complete its task.

A solutions architect must determine which permissions each Lambda function needs.

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

- A. Set up Amazon CodeGuru to profile the Lambda functions and search for AWS API call
- B. Create an inventory of the required API calls and resources for each Lambda function
- C. Create new IAM access policies for each Lambda function
- D. Review the new policies to ensure that they meet the company's business requirements.
- E. Turn on AWS CloudTrail logging for the AWS account
- F. Use AWS Identity and Access Management Access Analyzer to generate IAM access policies based on the activity recorded in the CloudTrail logs
- G. Review the generated policies to ensure that they meet the company's business requirements.
- H. Turn on AWS CloudTrail logging for the AWS account
- I. Create a script to parse the CloudTrail log, search for AWS API calls by Lambda execution role, and create a summary report
- J. Review the report
- K. Create IAM access policies that provide more restrictive permissions for each Lambda function.
- L. Turn on AWS CloudTrail logging for the AWS account
- M. Export the CloudTrail logs to Amazon S3. Use Amazon EMR to process the CloudTrail logs in Amazon S3 and produce a report of API calls and resources used by each execution role
- N. Create a new IAM access policy for each role
- O. Export the generated roles to an S3 bucket
- P. Review the generated policies to ensure that they meet the company's business requirements.

Answer: B

Explanation:

IAM Access Analyzer helps you identify the resources in your organization and accounts, such as Amazon S3 buckets or IAM roles, shared with an external entity. This lets you identify unintended access to your resources and data, which is a security risk. IAM Access Analyzer identifies resources shared with external principals by using logic-based reasoning to analyze the resource-based policies in your AWS environment.

<https://docs.aws.amazon.com/IAM/latest/UserGuide/what-is-access-analyzer.html>

NEW QUESTION 2

- (Exam Topic 2)

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

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

What should the solutions architect do to optimize the cost?

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

Answer: A

Explanation:

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

NEW QUESTION 3

- (Exam Topic 2)

A company has VPC flow logs enabled for its NAT gateway. The company is seeing Action = ACCEPT for inbound traffic that comes from public IP address 198.51.100.2 destined for a private Amazon EC2 instance.

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

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

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

L. Run a query to filter with the destination address set as "like 198.51.100.2" and the source address set as "like 203.0". Run the stats command to filter the sum of bytes transferred by the source address and the destination address.

Answer: D

Explanation:

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

NEW QUESTION 4

- (Exam Topic 2)

A company needs to optimize the cost of an AWS environment that contains multiple accounts in an organization in AWS Organizations. The company conducted cost optimization activities 3 years ago and purchased Amazon EC2 Standard Reserved Instances that recently expired.

The company needs EC2 instances for 3 more years. Additionally, the company has deployed a new serverless workload.

Which strategy will provide the company with the MOST cost savings?

- A. Purchase the same Reserved Instances for an additional 3-year term with All Upfront payment
- B. Purchase a 3-year Compute Savings Plan with All Upfront payment in the management account to cover any additional compute costs.
- C. Purchase a 1-year Compute Savings Plan with No Upfront payment in each member account
- D. Use the Savings Plans recommendations in the AWS Cost Management console to choose the Compute Savings Plan.
- E. Purchase a 3-year EC2 Instance Savings Plan with No Upfront payment in the management account to cover EC2 costs in each AWS Region
- F. Purchase a 3-year Compute Savings Plan with No Upfront payment in the management account to cover any additional compute costs.
- G. Purchase a 3-year EC2 Instance Savings Plan with All Upfront payment in each member account
- H. Use the Savings Plans recommendations in the AWS Cost Management console to choose the EC2 Instance Savings Plan.

Answer: A

Explanation:

The company should purchase the same Reserved Instances for an additional 3-year term with All Upfront payment. The company should purchase a 3-year Compute Savings Plan with All Upfront payment in the management account to cover any additional compute costs. This solution will provide the company with the most cost savings because Reserved Instances and Savings Plans are both pricing models that offer significant discounts compared to On-Demand pricing. Reserved Instances are commitments to use a specific instance type and size in a single Region for a one- or three-year term. You can choose between three payment options:

No Upfront, Partial Upfront, or All Upfront. The more you pay upfront, the greater the discount. Savings Plans are flexible pricing models that offer low prices on EC2 instances, Fargate, and Lambda usage, in exchange for a commitment to a consistent amount of usage (measured in \$/hour) for a one- or three-year term.

You can choose between two types of Savings Plans: Compute Savings Plans and EC2 Instance Savings Plans. Compute Savings Plans apply to any EC2 instance regardless of Region, instance family, operating system, or tenancy, including those that are part of EMR, ECS, or EKS clusters, or launched by Fargate or Lambda. EC2 Instance Savings Plans apply to a specific instance family within a Region and provide the most savings. By purchasing the same Reserved Instances for an additional 3-year term with All Upfront payment, the company can lock in the lowest possible price for its EC2 instances that run continuously for 3 years. By purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account, the company can benefit from additional discounts on any other compute usage across its member accounts.

The other options are not correct because:

➤ Purchasing a 1-year Compute Savings Plan with No Upfront payment in each member account would not provide as much cost savings as purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account. A 1-year term offers lower discounts than a 3-year term, and a No Upfront payment option offers lower discounts than an All Upfront payment option. Also, purchasing a Savings Plan in each member account would not allow the company to share the benefits of unused Savings Plan discounts across its organization.

➤ Purchasing a 3-year EC2 Instance Savings Plan with No Upfront payment in the management account to cover EC2 costs in each AWS Region would not provide as much cost savings as purchasing Reserved Instances for an additional 3-year term with All Upfront payment. An EC2 Instance Savings Plan offers lower discounts than Reserved Instances for the same instance family and Region. Also, a No Upfront payment option offers lower discounts than an All Upfront payment option.

➤ Purchasing a 3-year EC2 Instance Savings Plan with All Upfront payment in each member account would not provide as much flexibility or cost savings as purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account. An EC2 Instance Savings Plan applies only to a specific instance family within a Region and does not cover Fargate or Lambda usage. Also, purchasing a Savings Plan in each member account would not allow the company to share the benefits of unused Savings Plan discounts across its organization.

References:

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

NEW QUESTION 5

- (Exam Topic 2)

A solutions architect needs to improve an application that is hosted in the AWS Cloud. The application uses an Amazon Aurora MySQL DB instance that is experiencing overloaded connections. Most of the application's operations insert records into the database. The application currently stores credentials in a text-based configuration file.

The solutions architect needs to implement a solution so that the application can handle the current connection load. The solution must keep the credentials secure and must provide the ability to rotate the credentials automatically on a regular basis.

Which solution will meet these requirements?

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

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/rds-proxy.html>

NEW QUESTION 6

- (Exam Topic 2)

A company needs to build a disaster recovery (DR) solution for its ecommerce website. The web application is hosted on a fleet of t3.large Amazon EC2 instances and uses an Amazon RDS for MySQL DB instance. The EC2 instances are in an Auto Scaling group that extends across multiple Availability Zones. In the event of a disaster, the web application must fail over to the secondary environment with an RPO of 30 seconds and an RTO of 10 minutes. Which solution will meet these requirements MOST cost-effectively?

- A. Use infrastructure as code (IaC) to provision the new infrastructure in the DR Region
- B. Create a cross-Region read replica for the DB instance
- C. Set up a backup plan in AWS Backup to create cross-Region backups for the EC2 instances and the DB instance
- D. Create a cron expression to back up the EC2 instances and the DB instance every 30 seconds to the DR Region
- E. Recover the EC2 instances from the latest EC2 backup
- F. Use an Amazon Route 53 geolocation routing policy to automatically fail over to the DR Region in the event of a disaster.
- G. Use infrastructure as code (IaC) to provision the new infrastructure in the DR Region
- H. Create a cross-Region read replica for the DB instance
- I. Set up AWS Elastic Disaster Recovery to continuously replicate the EC2 instances to the DR Region
- J. Run the EC2 instances at the minimum capacity in the DR Region. Use an Amazon Route 53 failover routing policy to automatically fail over to the DR Region in the event of a disaster
- K. Increase the desired capacity of the Auto Scaling group.
- L. Set up a backup plan in AWS Backup to create cross-Region backups for the EC2 instances and the DB instance
- M. Create a cron expression to back up the EC2 instances and the DB instance every 30 seconds to the DR Region
- N. Use infrastructure as code (IaC) to provision the new infrastructure in the DR Region
- O. Manually restore the backed-up data on new instance
- P. Use an Amazon Route 53 simple routing policy to automatically fail over to the DR Region in the event of a disaster.
- Q. Use infrastructure as code (IaC) to provision the new infrastructure in the DR Region
- R. Create an Amazon Aurora global database
- S. Set up AWS Elastic Disaster Recovery to continuously replicate the EC2 instances to the DR Region
- T. Run the Auto Scaling group of EC2 instances at full capacity in the DR Region
- . Use an Amazon Route 53 failover routing policy to automatically fail over to the DR Region in the event of a disaster.

Answer: B

Explanation:

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

The other options are not correct because:

- Using AWS Backup to create cross-Region backups for the EC2 instances and the DB instance would not meet the RPO and RTO requirements. AWS Backup is a service that enables you to centralize and automate data protection across AWS services. You can use AWS Backup to back up your application data across AWS services in your account and across accounts. However, AWS Backup does not provide continuous replication or fast recovery; it creates backups at scheduled intervals and requires manual restoration. Creating backups every 30 seconds would also incur high costs and network bandwidth.
- Creating an Amazon API Gateway Data API service integration with Amazon Redshift would not help with disaster recovery. The Data API is a feature that enables you to query your Amazon Redshift cluster using HTTP requests, without needing a persistent connection or a SQL client. It is useful for building applications that interact with Amazon Redshift, but not for replicating or recovering data.
- Creating an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster would not help with disaster recovery. AWS Data Exchange is a service that makes it easy for AWS customers to exchange data in the cloud. You can use AWS Data Exchange to subscribe to a diverse selection of third-party data products or offer your own data products to other AWS customers. A datashare is a feature that enables you to share live and secure access to your Amazon Redshift data across your accounts or with third parties without copying or moving the underlying data. It is useful for sharing query results and views with other users, but not for replicating or recovering data.

References:

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

NEW QUESTION 7

- (Exam Topic 2)

A company is building a call center by using Amazon Connect. The company's operations team is defining a disaster recovery (DR) strategy across AWS Regions. The contact center has dozens of contact flows, hundreds of users, and dozens of claimed phone numbers. Which solution will provide DR with the LOWEST RTO?

- A. Create an AWS Lambda function to check the availability of the Amazon Connect instance and to send a notification to the operations team in case of

unavailabilit

- B. Create an Amazon EventBridge rule to invoke the Lambda function every 5 minute
- C. After notification, instruct the operations team to use theAWS Management Console to provision a new Amazon Connect instance in a second Regio
- D. Deploy the contact flows, users, and claimed phone numbers by using an AWS CloudFormation template.
- E. Provision a new Amazon Connect instance with all existing users in a second Regio
- F. Create an AWS Lambda function to check the availability of the Amazon Connect instanc
- G. Create an Amazon EventBridge rule to invoke the Lambda function every 5 minute
- H. In the event of an issue, configure the Lambda function to deploy an AWS CloudFormation template that provisions contact flows and claimed numbers in the second Region.
- I. Provision a new Amazon Connect instance with all existing contact flows and claimed phone numbers in a second Regio
- J. Create an Amazon Route 53 health check for the URL of the Amazon Connect instanc
- K. Create an Amazon CloudWatch alarm for failed health check
- L. Create an AWS Lambda function to deploy an AWS CloudFormation template that provisions all user
- M. Configure the alarm to invoke the Lambda function.
- N. Provision a new Amazon Connect instance with all existing users and contact flows in a second Region.Create an Amazon Route 53 health check for the URL of the Amazon Connect instanc
- O. Create an Amazon CloudWatch alarm for failed health check
- P. Create an AWS Lambda function to deploy an AWS CloudFormation template that provisions claimed phone number
- Q. Configure the alarm to invoke the Lambda function.

Answer: D

Explanation:

Option D provisions a new Amazon Connect instance with all existing users and contact flows in a second Region. It also sets up an Amazon Route 53 health check for the URL of the Amazon Connect instance, an Amazon CloudWatch alarm for failed health checks, and an AWS Lambda function to deploy an AWS CloudFormation template that provisions claimed phone numbers. This option allows for the fastest recovery time because all the necessary components are already provisioned and ready to go in the second Region. In the event of a disaster, the failed health check will trigger the AWS Lambda function to deploy the CloudFormation template to provision the claimed phone numbers, which is the only missing component.

NEW QUESTION 8

- (Exam Topic 2)

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

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

Which solution will meet these requirements MOST cost-effectively?

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

Answer: C

Explanation:

Generally, unstructured data should be converted structured data before querying them. AWS Glue can do that.

<https://docs.aws.amazon.com/glue/latest/dg/schema-relationalize.html> <https://docs.aws.amazon.com/athena/latest/ug/glue-athena.html>

NEW QUESTION 9

- (Exam Topic 2)

A solutions architect at a large company needs to set up network security tor outbound traffic to the internet from all AWS accounts within an organization in AWS Organizations. The organization has more than 100 AWS accounts, and the accounts route to each other by using a centralized AWS Transit Gateway. Each account has both an internet gateway and a NAT gateway tor outbound traffic to the internet The company deploys resources only into a single AWS Region.

The company needs the ability to add centrally managed rule-based filtering on all outbound traffic to the internet for all AWS accounts in the organization. The peak load of outbound traffic will not exceed 25 Gbps in each Availability Zone.

Which solution meets these requirements?

- A. Create a new VPC for outbound traffic to the interne
- B. Connect the existing transit gateway to the new VP
- C. Configure a new NAT gatewa
- D. Create an Auto Scaling group of Amazon EC2 instances that run an open-source internet proxy for rule-based filtering across all Availability Zones in the Regio
- E. Modify all default routes to point to the proxy's Auto Scaling group.
- F. Create a new VPC for outbound traffic to the interne
- G. Connect the existing transit gateway to the new VP
- H. Configure a new NAT gatewa
- I. Use an AWSNetwork Firewall firewall for rule-based filterin
- J. Create Network Firewall endpoints in each Availability Zon
- K. Modify all default routes to point to the Network Firewall endpoints.
- L. Create an AWS Network Firewall firewall for rule-based filtering in each AWS accoun
- M. Modify all default routes to point to the Network Firewall firewalls in each account.
- N. In each AWS account, create an Auto Scaling group of network-optimized Amazon EC2 instances that run an open-source internet proxy for rule-based filterin
- O. Modify all default routes to point to the proxy's Auto Scaling group.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/networking-and-content-delivery/deployment-models-for-aws-network-firewall/>

NEW QUESTION 10

- (Exam Topic 2)

A large company runs workloads in VPCs that are deployed across hundreds of AWS accounts. Each VPC consists to public subnets and private subnets that span across multiple Availability Zones. NAT gateways are deployed in the public subnets and allow outbound connectivity to the internet from the private subnets. A solutions architect is working on a hub-and-spoke design. All private subnets in the spoke VPCs must route traffic to the internet through an egress VPC. The solutions architect already has deployed a NAT gateway in an egress VPC in a central AWS account.

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

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

Answer: B

Explanation:

<https://d1.awsstatic.com/architecture-diagrams/ArchitectureDiagrams/NAT-gateway-centralized-egress-ra.pdf?d>

NEW QUESTION 10

- (Exam Topic 2)

A company manufactures smart vehicles. The company uses a custom application to collect vehicle data. The vehicles use the MQTT protocol to connect to the application.

The company processes the data in 5-minute intervals. The company then copies vehicle telematics data to on-premises storage. Custom applications analyze this data to detect anomalies.

The number of vehicles that send data grows constantly. Newer vehicles generate high volumes of data. The on-premises storage solution is not able to scale for peak traffic, which results in data loss. The company must modernize the solution and migrate the solution to AWS to resolve the scaling challenges.

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

- A. Use AWS IoT Greengrass to send the vehicle data to Amazon Managed Streaming for Apache Kafka (Amazon MSK). Create an Apache Kafka application to store the data in Amazon S3. Use a pretrained model in Amazon SageMaker to detect anomalies.
- B. Use AWS IoT Core to receive the vehicle data
- C. Configure rules to route data to an Amazon Kinesis Data Firehose delivery stream that stores the data in Amazon S3. Create an Amazon Kinesis Data Analytics application that reads from the delivery stream to detect anomalies.
- D. Use AWS IoT FleetWise to collect the vehicle data
- E. Send the data to an Amazon Kinesis data stream. Use an Amazon Kinesis Data Firehose delivery stream to store the data in Amazon S3. Use the built-in machine learning transforms in AWS Glue to detect anomalies.
- F. Use Amazon MQ for RabbitMQ to collect the vehicle data
- G. Send the data to an Amazon Kinesis Data Firehose delivery stream to store the data in Amazon S3. Use Amazon Lookout for Metrics to detect anomalies.

Answer: B

Explanation:

Using AWS IoT Core to receive the vehicle data will enable connecting the smart vehicles to the cloud using the MQTT protocol¹. AWS IoT Core is a platform that enables you to connect devices to AWS Services and other devices, secure data and interactions, process and act upon device data, and enable applications to interact with devices even when they are offline². Configuring rules to route data to an Amazon Kinesis Data Firehose delivery stream that stores the data in Amazon S3 will enable processing and storing the vehicle data in a scalable and reliable way³. Amazon Kinesis Data Firehose is a fully managed service that delivers real-time streaming data to destinations such as Amazon S3. Creating an Amazon Kinesis Data Analytics application that reads from the delivery stream to detect anomalies will enable analyzing the vehicle data using SQL queries or Apache Flink applications. Amazon Kinesis Data Analytics is a fully managed service that enables you to process and analyze streaming data using SQL or Java.

NEW QUESTION 12

- (Exam Topic 2)

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 17

- (Exam Topic 2)

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

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 22

- (Exam Topic 2)

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

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

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

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 24

- (Exam Topic 2)

A company is designing an AWS Organizations structure. The company wants to standardize a process to apply tags across the entire organization. The company will require tags with specific values when a user creates a new resource. Each of the company's OUs will have unique tag values.

Which solution will meet these requirements?

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

- J. Use an SCP to deny the creation of resources that do not have the required tag
- K. Define the list of tags. Attach the SCP to the OUs

Answer: A

Explanation:

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

NEW QUESTION 26

- (Exam Topic 2)

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 31

- (Exam Topic 2)

A company needs to establish a connection from its on-premises data center to AWS. The company needs to connect all of its VPCs that are located in different AWS Regions with transitive routing capabilities between VPC networks. The company also must reduce network outbound traffic costs, increase bandwidth throughput, and provide a consistent network experience for end users.

Which solution will meet these requirements?

- A. Create an AWS Site-to-Site VPN connection between the on-premises data center and a new central VP
- B. Create VPC peering connections that initiate from the central VPC to all other VPCs.
- C. Create an AWS Direct Connect connection between the on-premises data center and AW
- D. Provision a transit VIF, and connect it to a Direct Connect gatewa
- E. Connect the Direct Connect gateway to all the other VPCs by using a transit gateway in each Region.
- F. Create an AWS Site-to-Site VPN connection between the on-premises data center and a new central VP
- G. Use a transit gateway with dynamic routin
- H. Connect the transit gateway to all other VPCs.
- I. Create an AWS Direct Connect connection between the on-premises data center and AWS Establish an AWS Site-to-Site VPN connection between all VPCs in each Regio
- J. Create VPC peering connections that initiate from the central VPC to all other VPCs.

Answer: B

Explanation:

Transit GW + Direct Connect GW + Transit VIF + enabled SiteLink if two different DX locations <https://aws.amazon.com/blogs/networking-and-content-delivery/introducing-aws-direct-connect-sitelink/>

NEW QUESTION 33

- (Exam Topic 2)

A company has an application that runs as a ReplicaSet of multiple pods in an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The EKS cluster has nodes in multiple Availability Zones. The application generates many small files that must be accessible across all running instances of the application. The company needs to back up the files and retain the backups for 1 year.

Which solution will meet these requirements while providing the FASTEST storage performance?

- A. Create an Amazon Elastic File System (Amazon EFS) file system and a mount target for each subnet that contains nodes in the EKS cluste
- B. Configure the ReplicaSet to mount the file syste
- C. Direct the application to store files in the file syste
- D. Configure AWS Backup to back up and retain copies of the data for 1 year.
- E. Create an Amazon Elastic Block Store (Amazon EBS) volum
- F. Enable the EBS Multi-Attach feature. Configure the ReplicaSet to mount the EBS volum
- G. Direct the application to store files in the EBS volum
- H. Configure AWS Backup to back up and retain copies of the data for 1 year.
- I. Create an Amazon S3 bucke
- J. Configure the ReplicaSet to mount the S3 bucke
- K. Direct the application to store files in the S3 bucke
- L. Configure S3 Versioning to retain copies of the dat
- M. Configure an S3 Lifecycle policy to delete objects after 1 year.
- N. Configure the ReplicaSet to use the storage available on each of the running application pods to store the files locall
- O. Use a third-party tool to back up the EKS cluster for 1 year.

Answer: A

Explanation:

In the past, EBS can be attached only to one ec2 instance but not anymore but there are limitations like - it works only on io1/io2 instance types and many others as described here. <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volumes-multi.html> EFS has shareable storage

In terms of performance, Amazon EFS is optimized for workloads that require high levels of aggregate throughput and IOPS, whereas EBS is optimized for low-latency, random access I/O operations. Amazon EFS is designed to scale throughput and capacity automatically as your storage needs grow, while EBS volumes can be resized on demand.

NEW QUESTION 35

- (Exam Topic 2)

A company has built a high performance computing (HPC) cluster in AWS for a tightly coupled workload that generates a large number of shared files stored in Amazon EFS. The cluster was performing well when the number of Amazon EC2 instances in the cluster was 100. However, when the company increased the cluster size to 1,000 EC2 instances, overall performance was well below expectations.

Which collection of design choices should a solutions architect make to achieve the maximum performance from the HPC cluster? (Select THREE.)

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

Answer: ACF

Explanation:

* A. High performance computing (HPC) workload cluster should be in a single AZ.

* C. Elastic Fabric Adapter (EFA) is a network device that you can attach to your Amazon EC2 instances to accelerate High Performance Computing (HPC)

* F. Amazon FSx for Lustre - Use it for workloads where speed matters, such as machine learning, high performance computing (HPC), video processing, and financial modeling.

Cluster – packs instances close together inside an Availability Zone. This strategy enables workloads to achieve the low-latency network performance necessary for tightly-coupled node-to-node communication that is typical of HPC applications.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>

NEW QUESTION 37

- (Exam Topic 2)

A telecommunications company is running an application on AWS. The company has set up an AWS Direct Connect connection between the company's on-premises data center and AWS. The company deployed the application on Amazon EC2 instances in multiple Availability Zones behind an internal Application Load Balancer (ALB). The company's clients connect from the on-premises network by using HTTPS. The TLS terminates in the ALB. The company has multiple target groups and uses path-based routing to forward requests based on the URL path.

The company is planning to deploy an on-premises firewall appliance with an allow list that is based on IP address. A solutions architect must develop a solution to allow traffic flow to AWS from the on-premises network so that the clients can continue to access the application.

Which solution will meet these requirements?

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

Answer: B

Explanation:

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

NEW QUESTION 41

- (Exam Topic 2)

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

Which solution will meet these requirements?

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

Answer: A

Explanation:

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

NEW QUESTION 45

- (Exam Topic 2)

A company has five development teams that have each created five AWS accounts to develop and host applications. To track spending, the development teams log in to each account every month, record the current cost from the AWS Billing and Cost Management console, and provide the information to the company's finance team.

The company has strict compliance requirements and needs to ensure that resources are created only in AWS Regions in the United States. However, some resources have been created in other Regions.

A solutions architect needs to implement a solution that gives the finance team the ability to track and consolidate expenditures for all the accounts. The solution also must ensure that the company can create resources only in Regions in the United States.

Which combination of steps will meet these requirements in the MOST operationally efficient way? (Select THREE.)

- A. Create a new account to serve as a management account
- B. Create an Amazon S3 bucket for the finance team. Use AWS Cost and Usage Reports to create monthly reports and to store the data in the finance team's S3 bucket.
- C. Create a new account to serve as a management account
- D. Deploy an organization in AWS Organizations with all features enabled
- E. Invite all the existing accounts to the organization
- F. Ensure that each account accepts the invitation.
- G. Create an OU that includes all the development team
- H. Create an SCP that allows the creation of resources only in Regions that are in the United States
- I. Apply the SCP to the OU.
- J. Create an OU that includes all the development team
- K. Create an SCP that denies the creation of resources in Regions that are outside the United States
- L. Apply the SCP to the OU.
- M. Create an IAM role in the management account. Attach a policy that includes permissions to view the Billing and Cost Management console
- N. Allow the finance team users to assume the role
- O. Use AWS Cost Explorer and the Billing and Cost Management console to analyze cost.
- P. Create an IAM role in each AWS account
- Q. Attach a policy that includes permissions to view the Billing and Cost Management console
- R. Allow the finance team users to assume the role.

Answer: BCE

Explanation:

AWS Organizations is a service that enables you to consolidate multiple AWS accounts into an organization that you create and centrally manage. By creating a management account and inviting all the existing accounts to join the organization, the solutions architect can track and consolidate expenditures for all the accounts using AWS Cost Management tools such as AWS Cost Explorer and AWS Budgets. An organizational unit (OU) is a group of accounts within an organization that can be used to apply policies and simplify management. A service control policy (SCP) is a type of policy that you can use to manage permissions in your organization. By creating an OU that includes all the development teams and applying an SCP that allows the creation of resources only in Regions that are in the United States, the solutions architect can ensure that the company meets its compliance requirements and avoids unwanted charges from other Regions. An IAM role is an identity with permission policies that determine what the identity can and cannot do in AWS. By creating an IAM role in the management account and allowing the finance team users to assume it, the solutions architect can give them access to view the Billing and Cost Management console without sharing credentials or creating additional users. References:

- > https://docs.aws.amazon.com/organizations/latest/userguide/orgs_introduction.html
- > https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp.html
- > https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html
- > <https://docs.aws.amazon.com/aws-cost-management/latest/userguide/what-is-costmanagement.html>

NEW QUESTION 47

- (Exam Topic 2)

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

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 52

- (Exam Topic 1)

A company has deployed an application on AWS Elastic Beanstalk. The application uses Amazon Aurora for the database layer. An Amazon CloudFront distribution serves web requests and includes the Elastic Beanstalk domain name as the origin server. The distribution is configured with an alternate domain name that visitors use when they access the application.

Each week, the company takes the application out of service for routine maintenance. During the time that the application is unavailable, the company wants visitors to receive an informational message instead of a CloudFront error message.

A solutions architect creates an Amazon S3 bucket as the first step in the process.

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

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

Answer: ACD

Explanation:

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

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

The other options are not correct because:

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

References:

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

NEW QUESTION 53

- (Exam Topic 1)

A company is developing a new service that will be accessed using TCP on a static port A solutions architect must ensure that the service is highly available, has redundancy across Availability Zones, and is accessible using the DNS name myservice.com, which is publicly accessible The service must use fixed address assignments so other companies can add the addresses to their allow lists.

Assuming that resources are deployed in multiple Availability Zones in a single Region, which solution will meet these requirements?

- A. Create Amazon EC2 instances with an Elastic IP address for each instance Create a Network Load Balancer (NLB) and expose the static TCP port Register EC2instances with the NLB Create a new name server record set named my service com, and assign the Elastic IP addresses of the EC2 instances to the record set Provide the Elastic IP addresses of the EC2 instances to the other companies to add to their allow lists
- B. Create an Amazon ECS cluster and a service definition for the application Create and assign public IP addresses for the ECS cluster Create a Network Load Balancer (NLB) and expose the TCP port Create atarget group and assign the ECS cluster name to the NLB Create a new A record set named my service com and assign the public IP addresses of the ECS cluster to the record set Provide the public IP addresses of the ECS cluster to the other companies to add to their allow lists
- C. Create Amazon EC2 instances for the service Create one Elastic IP address for each Availability Zone Create a Network Load Balancer (NLB) and expose the assigned TCP port Assign the Elastic IP addresses to the NLB for each Availability Zone Create a target group and register the EC2 instances with the NLB Create a new A (alias) record set named my service com, and assign the NLB DNS name to the record set.
- D. Create an Amazon ECS cluster and a service definition for the application Create and assign public IP address for each host in the cluster Create an Application Load Balancer (ALB) and expose the static TCP port Create a target group and assign the ECS service definition name to the ALB Create a new CNAME record set and associate the public IP addresses to the record set Provide the Elastic IP addresses of the Amazon EC2 instances to the other companies to add to their allow lists

Answer: C

Explanation:

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-to-elb-load-balancer.html>

Create a Network Load Balancer (NLB) and expose the assigned TCP port. Assign the Elastic IP addresses to the NLB for each Availability Zone. Create a target group and register the EC2 instances with the NLB. Create a new A (alias) record set named my.service.com, and assign the NLB DNS name to the record set. As it uses the NLB as the resource in the A-record, traffic will be routed through the NLB, and it will automatically route the traffic to the healthy instances based on the health checks and also it provides the fixed address assignments as the other companies can add the NLB's Elastic IP addresses to their allow lists.

NEW QUESTION 56

- (Exam Topic 1)

A company is developing a new serverless API by using Amazon API Gateway and AWS Lambda. The company integrated the Lambda functions with API Gateway to use several shared libraries and custom classes.

A solutions architect needs to simplify the deployment of the solution and optimize for code reuse. Which solution will meet these requirements?

- A. Deploy the shared libraries and custom classes into a Docker imag
- B. Store the image in an S3 bucket.Create a Lambda layer that uses the Docker image as the sourc
- C. Deploy the API's Lambda functions as Zip package
- D. Configure the packages to use the Lambda layer.

- E. Deploy the shared libraries and custom classes to a Docker image
- F. Upload the image to Amazon Elastic Container Registry (Amazon ECR). Create a Lambda layer that uses the Docker image as the source
- G. Deploy the API's Lambda functions as Zip package
- H. Configure the packages to use the Lambda layer.
- I. Deploy the shared libraries and custom classes to a Docker container in Amazon Elastic Container Service (Amazon ECS) by using the AWS Fargate launch type
- J. Deploy the API's Lambda functions as Zip package
- K. Configure the packages to use the deployed container as a Lambda layer.
- L. Deploy the shared libraries, custom classes, and code for the API's Lambda functions to a Docker image
- M. Upload the image to Amazon Elastic Container Registry (Amazon ECR). Configure the API's Lambda functions to use the Docker image as the deployment package.

Answer: B

Explanation:

Deploying the shared libraries and custom classes to a Docker image and uploading the image to Amazon Elastic Container Registry (Amazon ECR) and creating a Lambda layer that uses the Docker image as the source. Then, deploying the API's Lambda functions as Zip packages and configuring the packages to use the Lambda layer would meet the requirements for simplifying the deployment and optimizing for code reuse.

A Lambda layer is a distribution mechanism for libraries, custom runtimes, and other function dependencies. It allows you to manage your in-development function code separately from your dependencies, this way you can easily update your dependencies without having to update your entire function code.

By deploying the shared libraries and custom classes to a Docker image and uploading the image to Amazon Elastic Container Registry (ECR), it makes it easy to manage and version the dependencies. This way, the company can use the same version of the dependencies across different Lambda functions.

By creating a Lambda layer that uses the Docker image as the source, the company can configure the API's Lambda functions to use the layer, reducing the need to include the dependencies in each function package, and making it easy to update the dependencies across all functions at once.

Reference:

AWS Lambda Layers documentation: <https://docs.aws.amazon.com/lambda/latest/dg/configuration-layers.html>

AWS Elastic Container Registry (ECR) documentation: <https://aws.amazon.com/ecr/> Building Lambda Layers with Docker documentation:

<https://aws.amazon.com/blogs/compute/building-lambda-layers-with-docker/>

NEW QUESTION 57

- (Exam Topic 1)

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

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

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

Answer: C

Explanation:

Set up an Amazon Aurora MySQL database. Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora. Create an Aurora Replica for the Aurora MySQL database, and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as AWS Lambda functions behind an Application Load Balancer (ALB), and use Amazon RDS Proxy to write to the Aurora MySQL database. When the databases are synced, point the collector DNS record to the ALB. Disable the AWS DMS sync task after the cutover from on premises to AWS.

Amazon RDS Proxy allows applications to pool and share connections established with the database, improving database efficiency and application scalability. With RDS Proxy, failover times for Aurora and RDS databases are reduced by up to 66%.

NEW QUESTION 60

- (Exam Topic 1)

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

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

- J. Enable AWS Cost and Usage Reports in the organization's master account and configure reports grouped by application, environment, and owner.
 K. Create an AWS Lambda function that processes AWS Cost and Usage Reports, sends budget alerts, and sends monthly reports to each business unit's email list.

Answer: B

Explanation:

Configure AWS Budgets in the organization's master account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in the organization's master account to create monthly reports for each business unit.
<https://aws.amazon.com/about-aws/whats-new/2019/07/introducing-aws-budgets-reports/#:~:text=AWS%20Bud>

NEW QUESTION 62

- (Exam Topic 1)

A company has migrated its forms-processing application to AWS. When users interact with the application, they upload scanned forms as files through a web application. A database stores user metadata and references to files that are stored in Amazon S3. The web application runs on Amazon EC2 instances and an Amazon RDS for PostgreSQL database.

When forms are uploaded, the application sends notifications to a team through Amazon Simple Notification Service (Amazon SNS). A team member then logs in and processes each form. The team member performs data validation on the form and extracts relevant data before entering the information into another system that uses an API.

A solutions architect needs to automate the manual processing of the forms. The solution must provide accurate form extraction, minimize time to market, and minimize long-term operational overhead.

Which solution will meet these requirements?

- A. Develop custom libraries to perform optical character recognition (OCR) on the form
- B. Deploy the libraries to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster as an application tier
- C. Use this tier to process the forms when forms are uploaded
- D. Store the output in Amazon S3. Parse this output by extracting the data into an Amazon DynamoDB table
- E. Submit the data to the target system's API
- F. Host the new application tier on EC2 instances.
- G. Extend the system with an application tier that uses AWS Step Functions and AWS Lambda
- H. Configure this tier to use artificial intelligence and machine learning (AI/ML) models that are trained and hosted on an EC2 instance to perform optical character recognition (OCR) on the forms when forms are uploaded
- I. Store the output in Amazon S3. Parse this output by extracting the data that is required within the application tier
- J. Submit the data to the target system's API.
- K. Host a new application tier on EC2 instance
- L. Use this tier to call endpoints that host artificial intelligence and machine learning (AI/ML) models that are trained and hosted in Amazon SageMaker to perform optical character recognition (OCR) on the form
- M. Store the output in Amazon ElastiCache
- N. Parse this output by extracting the data that is required within the application tier
- O. Submit the data to the target system's API.
- P. Extend the system with an application tier that uses AWS Step Functions and AWS Lambda
- Q. Configure this tier to use Amazon Textract and Amazon Comprehend to perform optical character recognition (OCR) on the forms when forms are uploaded
- R. Store the output in Amazon S3. Parse this output by extracting the data that is required within the application tier
- S. Submit the data to the target system's API.

Answer: D

Explanation:

Extend the system with an application tier that uses AWS Step Functions and AWS Lambda. Configure this tier to use Amazon Textract and Amazon Comprehend to perform optical character recognition (OCR) on the forms when forms are uploaded. Store the output in Amazon S3. Parse this output by extracting the data that is required within the application tier. Submit the data to the target system's API. This solution meets the requirements of accurate form extraction, minimal time to market, and minimal long-term operational overhead. Amazon Textract and Amazon Comprehend are fully managed and serverless services that can perform OCR and extract relevant data from the forms, which eliminates the need to develop custom libraries or train and host models. Using AWS Step Functions and Lambda allows for easy automation of the process and the ability to scale as needed.

NEW QUESTION 67

- (Exam Topic 1)

A company recently deployed an application on AWS. The application uses Amazon DynamoDB. The company measured the application load and configured the RCUs and WCUs on the DynamoDB table to match the expected peak load. The peak load occurs once a week for a 4-hour period and is double the average load. The application load is close to the average load for the rest of the week. The access pattern includes many more writes to the table than reads of the table.

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

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

Answer: D

Explanation:

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

NEW QUESTION 70

- (Exam Topic 1)

A company is subject to regulatory audits of its financial information. External auditors who use a single AWS account need access to the company's AWS account. A solutions architect must provide the auditors with secure, read-only access to the company's AWS account. The solution must comply with AWS security best practices.

Which solution will meet these requirements?

- A. In the company's AWS account, create resource policies for all resources in the account to grant access to the auditors' AWS account
- B. Assign a unique external ID to the resource policy.
- C. In the company's AWS account create an IAM role that trusts the auditors' AWS account Create an IAM policy that has the required permission
- D. Attach the policy to the rol
- E. Assign a unique external ID to the role's trust policy.
- F. In the company's AWS account, create an IAM use
- G. Attach the required IAM policies to the IAM user.Create API access keys for the IAM use
- H. Share the access keys with the auditors.
- I. In the company's AWS account, create an IAM group that has the required permissions Create an IAM user in the company s account for each audito
- J. Add the IAM users to the IAM group.

Answer: B

Explanation:

This solution will allow the external auditors to have read-only access to the company's AWS account while being compliant with AWS security best practices. By creating an IAM role, which is a secure and flexible way of granting access to AWS resources, and trusting the auditors' AWS account, the company can ensure that the auditors only have the permissions that are required for their role and nothing more. Assigning a unique external ID to the role's trust policy, it will ensure that only the auditors' AWS account can assume the role.

Reference:

AWS IAM Roles documentation: <https://aws.amazon.com/iam/features/roles/> AWS IAM Best practices: <https://aws.amazon.com/iam/security-best-practices/>

NEW QUESTION 72

- (Exam Topic 1)

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 h 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 rol
- D. The partner company should then use the IAM rote's Amazon Resource Name (ARN) when requesting access to perform the required tasks
- E. The customer should create an IAM rote and assign the required permissions to the IAM rot
- F. The partner company should then use the IAM rote's Amazon Resource Name (ARN). Including the external ID in the IAM role's trust pokey, when requesting access to perform the required tasks

Answer: C

Explanation:

<https://docs.aws.amazon.com/IAM/latest/UserGuide/confused-deputy.html>

This is the most secure way to allow org1 to access resources in org2 because it allows for least privilege security access. The customer should create an IAM role and assign the required permissions to the IAM role. The partner company should then use the IAM role's Amazon Resource Name (ARN) and include the external ID in the IAM role's trust policy when requesting access to perform the required tasks. This ensures that the partner company can only access the resources that it needs and only from the specific customer account.

NEW QUESTION 76

- (Exam Topic 1)

A company developed a pilot application by using AWS Elastic Beanstalk and Java. To save costs during development, the company's development team deployed the application into a single-instance environment. Recent tests indicate that the application consumes more CPU than expected. CPU utilization is regularly greater than 85%, which causes some performance bottlenecks.

A solutions architect must mitigate the performance issues before the company launches the application to production.

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

- A. Create a new Elastic Beanstalk applicatio
- B. Select a load-balanced environment typ
- C. Select all Availability Zone
- D. Add a scale-out rule that will run if the maximum CPU utilization is over 85% for 5 minutes.
- E. Create a second Elastic Beanstalk environmen
- F. Apply the traffic-splitting deployment polic
- G. Specify a percentage of incoming traffic to direct to the new environment in the average CPU utilization is over 85% for 5 minutes.
- H. Modify the existing environment's capacity configuration to use a load-balanced environment type.Select all Availability Zone
- I. Add a scale-out rule that will run if the average CPU utilization is over 85% for 5 minutes.
- J. Select the Rebuild environment action with the load balancing option Select an Availability Zones Add a scale-out rule that will run if the sum CPU utilization is over 85% for 5 minutes.

Answer: C

Explanation:

This solution will meet the requirements with the least operational overhead because it allows the company to modify the existing environment's capacity configuration, so it becomes a load-balanced environment type. By selecting all availability zones, the company can ensure that the application is running in multiple availability zones, which can help to improve the availability and scalability of the application. The company can also add a scale-out rule that will run if the average CPU utilization is over 85% for 5 minutes, which can help to mitigate the performance issues. This solution does not require creating new Elastic Beanstalk environments or rebuilding the existing one, which reduces the operational overhead.

You can refer to the AWS Elastic Beanstalk documentation for more information on how to use this service: <https://aws.amazon.com/elasticbeanstalk/> You can refer to the AWS documentation for more information on how to use autoscaling: <https://aws.amazon.com/autoscaling/>

NEW QUESTION 77

- (Exam Topic 1)

A company has developed a web application. The company is hosting the application on a group of Amazon EC2 instances behind an Application Load Balancer. The company wants to improve the security posture of the application and plans to use AWS WAF web ACLs. The solution must not adversely affect legitimate traffic to the application.

How should a solutions architect configure the web ACLs to meet these requirements?

- A. Set the action of the web ACL rules to Count
- B. Enable AWS WAF logging Analyze the requests for false positives Modify the rules to avoid any false positive Over time change the action of the web ACL rules from Count to Block.
- C. Use only rate-based rules in the web ACL
- D. and set the throttle limit as high as possible Temporarily block all requests that exceed the limit
- E. Define nested rules to narrow the scope of the rate tracking.
- F. Set the action of the web ACL rules to Block
- G. Use only AWS managed rule groups in the web ACLs Evaluate the rule groups by using Amazon CloudWatch metrics with AWS WAF sampled requests or AWS WAF logs.
- H. Use only custom rule groups in the web ACL
- I. and set the action to Allow Enable AWS WAF logging Analyze the requests for false positives Modify the rules to avoid any false positive Over time, change the action of the web ACL rules from Allow to Block.

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/waf-analyze-count-action-rules/>

NEW QUESTION 79

- (Exam Topic 1)

A company has a latency-sensitive trading platform that uses Amazon DynamoDB as a storage backend. The company configured the DynamoDB table to use on-demand capacity mode. A solutions architect needs to design a solution to improve the performance of the trading platform. The new solution must ensure high availability for the trading platform.

Which solution will meet these requirements with the LEAST latency?

- A. Create a two-node DynamoDB Accelerator (DAX) cluster Configure an application to read and write data by using DAX.
- B. Create a three-node DynamoDB Accelerator (DAX) cluster
- C. Configure an application to read data by using DAX and to write data directly to the DynamoDB table.
- D. Create a three-node DynamoDB Accelerator (DAX) cluster
- E. Configure an application to read data directly from the DynamoDB table and to write data by using DAX.
- F. Create a single-node DynamoDB Accelerator (DAX) cluster
- G. Configure an application to read data by using DAX and to write data directly to the DynamoDB table.

Answer: B

Explanation:

A DAX cluster can be deployed with one or two nodes for development or test workloads. One- and two-node clusters are not fault-tolerant, and we don't recommend using fewer than three nodes for production use. If a one- or two-node cluster encounters software or hardware errors, the cluster can become unavailable or lose cached data. A DAX cluster can be deployed with one or two nodes for development or test workloads. One and two-node clusters are not fault-tolerant, and we don't recommend using fewer than three nodes for production use. If a one- or two-node cluster encounters software or hardware errors, the cluster can become unavailable or lose cached data.

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DAX.concepts.cluster.html>

NEW QUESTION 84

- (Exam Topic 1)

A company plans to refactor a monolithic application into a modern application designed for AWS. The CI/CD pipeline needs to be upgraded to support the modern design for the application with the following requirements

- It should allow changes to be released several times every hour.
- * It should be able to roll back the changes as quickly as possible Which design will meet these requirements?

- A. Deploy a CI-CD pipeline that incorporates AMIs to contain the application and their configurations Deploy the application by replacing Amazon EC2 instances
- B. Specify AWS Elastic Beanstalk to serve as a secondary environment as the deployment target for the CI/CD pipeline of the application
- C. To deploy swap the staging and production environment URLs.
- D. Use AWS Systems Manager to re-provision the infrastructure for each deployment Update the Amazon EC2 user data to pull the latest code artifact from Amazon S3 and use Amazon Route 53 weighted routing to point to the new environment
- E. Roll out application updates as part of an Auto Scaling event using prebuilt AMI
- F. Use new versions of the AMIs to add instances, and phase out all instances that use the previous AMI version with the configured termination policy during a deployment event.

Answer: B

Explanation:

It is the fastest when it comes to rollback and deploying changes every hour

NEW QUESTION 87

- (Exam Topic 1)

A finance company is running its business-critical application on current-generation Linux EC2 instances The application includes a self-managed MySQL database performing heavy I/O operations. The application is working fine to handle a moderate amount of traffic during the month. However, it slows down during the final three days of each month due to month-end reporting, even though the company is using Elastic Load Balancers and Auto Scaling within its infrastructure to meet the increased demand.

Which of the following actions would allow the database to handle the month-end load with the LEAST impact on performance?

- A. Pre-warming Elastic Load Balancers, using a bigger instance type, changing all Amazon EBS volumes to GP2 volumes.
- B. Performing a one-time migration of the database cluster to Amazon RD
- C. and creating several additional read replicas to handle the load during end of month
- D. Using Amazon CioudWatch with AWS Lambda to change the typ
- E. size, or IOPS of Amazon EBS volumes in the cluster based on a specific CloudWatch metric
- F. Replacing all existing Amazon EBS volumes with new PIOPS volumes that have the maximum available storage size and I/O per second by taking snapshots before the end of the month and reverting back afterwards.

Answer: B

Explanation:

In this scenario, the Amazon EC2 instances are in an Auto Scaling group already which means that the database read operations is the possible bottleneck especially during the month-end wherein the reports are generated. This can be solved by creating RDS read replicas.

NEW QUESTION 88

- (Exam Topic 1)

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 AH 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 evenin
- C. Configure the rule to invoke an AWS Lambda function that stops instances based on thetag-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 rule that runs every hou
- F. Configure the rule to invoke one AWS Lambda function that terminates or restores instances from their last backup based on the ta
- G. day, and time.

Answer: B

Explanation:

Creating an Amazon EventBridge rule that runs every business day in the evening to stop instances and another rule that runs every business day in the morning to start instances based on the tag will reduce costs with the least operational effort. This approach allows for instances to be stopped during non-business hours when they are not in use, reducing the costs associated with running them. It also allows for instances to be started again in the morning when the development and testing teams need to use them.

NEW QUESTION 89

- (Exam Topic 1)

A company uses an on-premises data analytics platform. The system is highly available in a fully redundant configuration across 12 servers in the company's data center.

The system runs scheduled jobs, both hourly and daily, in addition to one-time requests from users. Scheduled jobs can take between 20 minutes and 2 hours to finish running and have tight SLAs. The scheduled jobs account for 65% of the system usage. User jobs typically finish running in less than 5 minutes and have no SLA. The user jobs account for 35% of system usage. During system failures, scheduled jobs must continue to meet SLAs. However, user jobs can be delayed.

A solutions architect needs to move the system to Amazon EC2 instances and adopt a consumption-based model to reduce costs with no long-term commitments.

The solution must maintain high availability and must not affect the SLAs.

Which solution will meet these requirements MOST cost-effectively?

- A. Split the 12 instances across two Availability Zones in the chosen AWS Regio
- B. Run two instances in each Availability Zone as On-Demand Instances with Capacity Reservation
- C. Run four instances in each Availability Zone as Spot Instances.
- D. Split the 12 instances across three Availability Zones in the chosen AWS Regio
- E. In one of the Availability Zones, run all four instances as On-Demand Instances with Capacity Reservation
- F. Run the remaining instances as Spot Instances.
- G. Split the 12 instances across three Availability Zones in the chosen AWS Regio
- H. Run two instances in each Availability Zone as On-Demand Instances with a Savings Pla
- I. Run two instances in each Availability Zone as Spot Instances.
- J. Split the 12 instances across three Availability Zones in the chosen AWS Regio
- K. Run three instances in each Availability Zone as On-Demand Instances with Capacity Reservation
- L. Run one instance in each Availability Zone as a Spot Instance.

Answer: D

Explanation:

By splitting the 12 instances across three Availability Zones, the system can maintain high availability and availability of resources in case of a failure. Option D also uses a combination of On-Demand Instances with Capacity Reservations and Spot Instances, which allows for scheduled jobs to be run on the On-Demand instances with guaranteed capacity, while also taking advantage of the cost savings from Spot Instances for the user jobs which have lower SLA requirements.

NEW QUESTION 93

- (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 jo
- 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 templat
- G. Use the EBS volume as the shared storage for the duration of the jo
- 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 clas
- 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 jo
- L. Delete the file system when the job is complete.
- M. Migrate the data from the existing shared file system to an Amazon S3 bucke
- 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 jo
- O. Delete the file gateway when the job is complete.

Answer: A

Explanation:

<https://aws.amazon.com/blogs/storage/new-enhancements-for-moving-data-between-amazon-fsx-for-lustre-and>

NEW QUESTION 95

- (Exam Topic 1)

A company with global offices has a single 1 Gbps AWS Direct Connect connection to a single AWS Region. The company's on-premises network uses the connection to communicate with the company's resources in the AWS Cloud. The connection has a single private virtual interface that connects to a single VPC. A solutions architect must implement a solution that adds a redundant Direct Connect connection in the same Region. The solution also must provide connectivity to other Regions through the same pair of Direct Connect connections as the company expands into other Regions.

Which solution meets these requirements?

- A. Provision a Direct Connect gatewa
- B. Delete the existing private virtual interface from the existing connectio
- C. Create the second Direct Connect connectio
- D. Create a new private virtual interlace on each connection, and connect both private virtual interfaces to the Direct Connect gatewa
- E. Connect the Direct Connect gateway to the single VPC.
- F. Keep the existing private virtual interfac
- G. Create the second Direct Connect connectio
- H. Create a new private virtual interface on the new connection, and connect the new private virtual interface to the single VPC.
- I. Keep the existing private virtual interfac
- J. Create the second Direct Connect connectio
- K. Create a new public virtual interface on the new connection, and connect the new public virtual interface to the single VPC.
- L. Provision a transit gatewa
- M. Delete the existing private virtual interface from the existing connection. Create the second Direct Connect connectio
- N. Create a new private virtual interface on each connection, and connect both private virtual interfaces to the transit gatewa
- O. Associate the transit gateway with the single VPC.

Answer: A

Explanation:

A Direct Connect gateway is a globally available resource. You can create the Direct Connect gateway in any Region and access it from all other Regions. The following describe scenarios where you can use a Direct Connect gateway.

<https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-gateways-intro.html>

NEW QUESTION 97

- (Exam Topic 1)

A company built an application based on AWS Lambda deployed in an AWS CloudFormation stack. The last production release of the web application introduced an issue that resulted in an outage lasting several minutes. A solutions architect must adjust the deployment process to support a canary release.

Which solution will meet these requirements?

- A. Create an alias for every new deployed version of the Lambda functio
- B. Use the AWS CLI update-alias command with the routing-config parameter to distribute the load.
- C. Deploy the application into a new CloudFormation stac
- D. Use an Amazon Route 53 weighted routing policy to distribute the load.
- E. Create a version for every new deployed Lambda functio
- F. Use the AWS CLI update-function-configuration command with the routing-config parameter to distribute the load.
- G. Configure AWS CodeDeploy and use CodeDeployDefault.OneAtATime in the Deployment configuration to distribute the load.

Answer: A

Explanation:

<https://aws.amazon.com/blogs/compute/implementing-canary-deployments-of-aws-lambda-functions-with-alias->
<https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html>

NEW QUESTION 100

- (Exam Topic 1)

A company is using Amazon OpenSearch Service to analyze data. The company loads data into an OpenSearch Service cluster with 10 data nodes from an

Amazon S3 bucket that uses S3 Standard storage. The data resides in the cluster for 1 month for read-only analysis. After 1 month, the company deletes the index that contains the data from the cluster. For compliance purposes, the company must retain a copy of all input data. The company is concerned about ongoing costs and asks a solutions architect to recommend a new solution. Which solution will meet these requirements MOST cost-effectively?

- A. Replace all the data nodes with UltraWarm nodes to handle the expected capacity
- B. Transition the input data from S3 Standard to S3 Glacier Deep Archive when the company loads the data into the cluster.
- C. Reduce the number of data nodes in the cluster to 2. Add UltraWarm nodes to handle the expected capacity
- D. Configure the indexes to transition to UltraWarm when OpenSearch Service ingests the data
- E. Transition the input data to S3 Glacier Deep Archive after 1 month by using an S3 Lifecycle policy.
- F. Reduce the number of data nodes in the cluster to 2. Add UltraWarm nodes to handle the expected capacity
- G. Configure the indexes to transition to UltraWarm when OpenSearch Service ingests the data
- H. Add cold storage nodes to the cluster. Transition the indexes from UltraWarm to cold storage
- I. Delete the input data from the S3 bucket after 1 month by using an S3 Lifecycle policy.
- J. Reduce the number of data nodes in the cluster to 2. Add instance-backed data nodes to handle the expected capacity
- K. Transition the input data from S3 Standard to S3 Glacier Deep Archive when the company loads the data into the cluster.

Answer: B

Explanation:

By reducing the number of data nodes in the cluster to 2 and adding UltraWarm nodes to handle the expected capacity, the company can reduce the cost of running the cluster. Additionally, configuring the indexes to transition to UltraWarm when OpenSearch Service ingests the data will ensure that the data is stored in the most cost-effective manner. Finally, transitioning the input data to S3 Glacier Deep Archive after 1 month by using an S3 Lifecycle policy will ensure that the data is retained for compliance purposes, while also reducing the ongoing costs.

NEW QUESTION 103

- (Exam Topic 1)

A company uses a service to collect metadata from applications that the company hosts on premises. Consumer devices such as TVs and internet radios access the applications. Many older devices do not support certain HTTP headers and exhibit errors when these headers are present in responses. The company has configured an on-premises load balancer to remove the unsupported headers from responses sent to older devices, which the company identified by the User-Agent headers.

The company wants to migrate the service to AWS, adopt serverless technologies, and retain the ability to support the older devices. The company has already migrated the applications into a set of AWS Lambda functions.

Which solution will meet these requirements?

- A. Create an Amazon CloudFront distribution for the metadata service
- B. Create an Application Load Balancer (ALB). Configure the CloudFront distribution to forward requests to the ALB
- C. Configure the ALB to invoke the correct Lambda function for each type of request
- D. Create a CloudFront function to remove the problematic headers based on the value of the User-Agent header.
- E. Create an Amazon API Gateway REST API for the metadata service
- F. Configure API Gateway to invoke the correct Lambda function for each type of request
- G. Modify the default gateway responses to remove the problematic headers based on the value of the User-Agent header.
- H. Create an Amazon API Gateway HTTP API for the metadata service
- I. Configure API Gateway to invoke the correct Lambda function for each type of request
- J. Create a response mapping template to remove the problematic headers based on the value of the User-Agent header
- K. Associate the response data mapping with the HTTP API.
- L. Create an Amazon CloudFront distribution for the metadata service
- M. Create an Application Load Balancer (ALB). Configure the CloudFront distribution to forward requests to the ALB
- N. Configure the ALB to invoke the correct Lambda function for each type of request
- O. Create a Lambda@Edge function that will remove the problematic headers in response to viewer requests based on the value of the User-Agent header.

Answer: D

Explanation:

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

NEW QUESTION 104

- (Exam Topic 1)

A company with several AWS accounts is using AWS Organizations and service control policies (SCPs). An Administrator created the following SCP and has attached it to an organizational unit (OU) that contains AWS account 1111-1111-1111:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAllActions",
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*"
    },
    {
      "Sid": "DenyCloudTrail",
      "Effect": "Deny",
      "Action": "cloudtrail:*",
      "Resource": "*"
    }
  ]
}
```

Developers working in account 1111-1111-1111 complain that they cannot create Amazon S3 buckets. How should the Administrator address this problem?

- A. Add s3:CreateBucket with Allow effect to the SCP.
- B. Remove the account from the OU, and attach the SCP directly to account 1111-1111-1111.
- C. Instruct the Developers to add Amazon S3 permissions to their IAM entities.
- D. Remove the SCP from account 1111-1111-1111.

Answer: C

Explanation:

However A's explanation is incorrect - https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html

"SCPs are similar to AWS Identity and Access Management (IAM) permission policies and use almost the same syntax. However, an SCP never grants permissions."

SCPs alone are not sufficient to granting permissions to the accounts in your organization. No permissions are granted by an SCP. An SCP defines a guardrail, or sets limits, on the actions that the account's administrator can delegate to the IAM users and roles in the affected accounts. The administrator must still attach identity-based or resource-based policies to IAM users or roles, or to the resources in your accounts to actually grant permissions. The effective permissions are the logical intersection between what is allowed by the SCP and what is allowed by the IAM and resource-based policies.

NEW QUESTION 106

- (Exam Topic 1)

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 bucket
- 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: C

Explanation:

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/high_availability_origin_failover.html

NEW QUESTION 109

- (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>

NEW QUESTION 110

- (Exam Topic 1)

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

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 114

- (Exam Topic 1)

A company is running several workloads in a single AWS account. A new company policy states that engineers can provision only approved resources and that engineers must use AWS CloudFormation to provision these resources. A solutions architect needs to create a solution to enforce the new restriction on the IAM role that the engineers use for access.

What should the solutions architect do to create the solution?

- A. Upload AWS CloudFormation templates that contain approved resources to an Amazon S3 bucket. Update the IAM policy for the engineers' IAM role to only allow access to Amazon S3 and AWS CloudFormation
- B. Use AWS CloudFormation templates to provision resources.
- C. Update the IAM policy for the engineers' IAM role with permissions to only allow provisioning of approved resources and AWS CloudFormation
- D. Use AWS CloudFormation templates to create stacks with approved resources.
- E. Update the IAM policy for the engineers' IAM role with permissions to only allow AWS CloudFormation actions
- F. Create a new IAM policy with permission to provision approved resources, and assign the policy to a new IAM service role
- G. Assign the IAM service role to AWS CloudFormation during stack creation.
- H. Provision resources in AWS CloudFormation stack
- I. Update the IAM policy for the engineers' IAM role to only allow access to their own AWS CloudFormation stack.

Answer: B

Explanation:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/security-best-practices.html#use-iam-to-c>

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-iam-service-role.html>

NEW QUESTION 117

- (Exam Topic 1)

The company needs to determine which costs on the monthly AWS bill are attributable to each application or team. The company also must be able to create reports to compare costs from the last 12 months and to help forecast costs for the next 12 months. A solutions architect must recommend an AWS Billing and Cost Management solution that provides these cost reports.

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

- A. Activate the user-defined cost allocation tags that represent the application and the team.
- B. Activate the AWS generated cost allocation tags that represent the application and the team.
- C. Create a cost category for each application in Billing and Cost Management.
- D. Activate IAM access to Billing and Cost Management.
- E. Create a cost budget.
- F. Enable Cost Explorer.

Answer: ACF

Explanation:

<https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/manage-cost-categories.html> <https://aws.amazon.com/premiumsupport/knowledge-center/cost-explorer-analyze-spending-and-usage/> <https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/manage-cost-categories.html>

<https://docs.aws.amazon.com/cost-management/latest/userguide/ce-enable.html>

The best combination of actions to meet the company's requirements is Options A, C, and F.

Option A involves activating the user-defined cost allocation tags that represent the application and the team. This will allow the company to assign costs to different applications or teams, and will allow them to be tracked in the monthly AWS bill.

Option C involves creating a cost category for each application in Billing and Cost Management. This will allow the company to easily identify and compare costs across different applications and teams.

Option F involves enabling Cost Explorer. This will allow the company to view the costs of their AWS resources over the last 12 months and to create forecasts for

the next 12 months.

These recommendations are in line with the official Amazon Textbook and Resources for the AWS Certified Solutions Architect - Professional certification. In particular, the book states that "You can use cost allocation tags to group your costs by application, team, or other categories" (Source: https://d1.awsstatic.com/training-and-certification/docs-sa-pro/AWS_Certified_Solutions_Architect_Professiona). Additionally, the book states that "Cost Explorer enables you to view the costs of your AWS resources over the last 12 months and to create forecasts for the next 12 months" (Source: https://d1.awsstatic.com/training-and-certification/docs-sa-pro/AWS_Certified_Solutions_Architect_Professiona).

NEW QUESTION 118

- (Exam Topic 1)

A company has hundreds of AWS accounts. The company recently implemented a centralized internal process for purchasing new Reserved Instances and modifying existing Reserved Instances. This process requires all business units that want to purchase or modify Reserved Instances to submit requests to a dedicated team for procurement. Previously, business units directly purchased or modified Reserved Instances in their own respective AWS accounts autonomously.

A solutions architect needs to enforce the new process in the most secure way possible.

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

- A. Ensure that all AWS accounts are part of an organization in AWS Organizations with all features enabled.
- B. Use AWS Config to report on the attachment of an IAM policy that denies access to the `ec2:PurchaseReservedInstancesOffering` action and the `ec2:ModifyReservedInstances` action.
- C. In each AWS account, create an IAM policy that denies the `ec2:PurchaseReservedInstancesOffering` action and the `ec2:ModifyReservedInstances` action.
- D. Create an SCP that denies the `ec2:PurchaseReservedInstancesOffering` action and the `ec2:ModifyReservedInstances` action.
- E. Attach the SCP to each OU of the organization.
- F. Ensure that all AWS accounts are part of an organization in AWS Organizations that uses the consolidated billing feature.

Answer: AD

Explanation:

All features – The default feature set that is available to AWS Organizations. It includes all the functionality of consolidated billing, plus advanced features that give you more control over accounts in your organization. For example, when all features are enabled the management account of the organization has full control over what member accounts can do. The management account can apply SCPs to restrict the services and actions that users (including the root user) and roles in an account can access. https://docs.aws.amazon.com/organizations/latest/userguide/orgs_getting-started_concepts.html#feature-set

NEW QUESTION 123

- (Exam Topic 1)

A company has a legacy monolithic application that is critical to the company's business. The company hosts the application on an Amazon EC2 instance that runs Amazon Linux 2. The company's application team receives a directive from the legal department to back up the data from the instance's encrypted Amazon Elastic Block Store (Amazon EBS) volume to an Amazon S3 bucket. The application team does not have the administrative SSH key pair for the instance. The application must continue to serve the users.

Which solution will meet these requirements?

- A. Attach a role to the instance with permission to write to Amazon S3. Use the AWS Systems Manager Session Manager option to gain access to the instance and run commands to copy data into Amazon S3.
- B. Create an image of the instance with the reboot option turned on.
- C. Launch a new EC2 instance from the image.
- D. Attach a role to the new instance with permission to write to Amazon S3. Run a command to copy data into Amazon S3.
- E. Take a snapshot of the EBS volume by using Amazon Data Lifecycle Manager (Amazon DLM). Copy the data to Amazon S3.
- F. Create an image of the instance.
- G. Launch a new EC2 instance from the image.
- H. Attach a role to the new instance with permission to write to Amazon S3. Run a command to copy data into Amazon S3.

Answer: C

Explanation:

Taking a snapshot of the EBS volume using Amazon Data Lifecycle Manager (DLM) will meet the requirements because it allows you to create a backup of the volume without the need to access the instance or its SSH key pair. Additionally, DLM allows you to schedule the backups to occur at specific intervals and also enables you to copy the snapshots to an S3 bucket. This approach will not impact the running application as the backup is performed on the EBS volume level.

NEW QUESTION 124

- (Exam Topic 1)

A company used Amazon EC2 instances to deploy a web fleet to host a blog site. The EC2 instances are behind an Application Load Balancer (ALB) and are configured in an Auto Scaling group. The web application stores all blog content on an Amazon EFS volume.

The company recently added a feature for bloggers to add video to their posts, attracting 10 times the previous user traffic. At peak times of day, users report buffering and timeout issues while attempting to reach the site or watch videos.

Which is the MOST cost-efficient and scalable deployment that will resolve the issues for users?

- A. Reconfigure Amazon EFS to enable maximum I/O.
- B. Update the blog site to use instance store volumes for storage.
- C. Copy the site contents to the volumes at launch and to Amazon S3 at shutdown.
- D. Configure an Amazon CloudFront distribution.
- E. Point the distribution to an S3 bucket, and migrate the videos from EFS to Amazon S3.
- F. Set up an Amazon CloudFront distribution for all site contents, and point the distribution at the ALB.

Answer: C

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-https-connection-fails/> Using an Amazon S3 bucket

Using a MediaStore container or a MediaPackage channel Using an Application Load Balancer

Using a Lambda function URL

Using Amazon EC2 (or another custom origin)

Using CloudFront origin groups <https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/restrict-access-to-load-balancer.html>

NEW QUESTION 127

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AWS-Certified-Solutions-Architect-Professional Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AWS-Certified-Solutions-Architect-Professional Product From:

<https://www.2passeasy.com/dumps/AWS-Certified-Solutions-Architect-Professional/>

Money Back Guarantee

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

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