

Exam Questions SAP-C01

AWS Certified Solutions Architect- Professional

<https://www.2passeasy.com/dumps/SAP-C01/>



NEW QUESTION 1

A company is currently running a production workload on AWS that is very I/O intensive. Its workload consists of a single tier with 10 c4.8xlarge instances, each with 2 TB gp2 volumes. The number of processing jobs has recently increased, and latency has increased as well. The team realizes that they are constrained on the IOPS. For the application to perform efficiently, they need to increase the IOPS by 3,000 for each of the instances. Which of the following designs will meet the performance goal MOST cost effectively?

- A. Change the type of Amazon EBS volume from gp2 to io1 and set provisioned IOPS to 9,000.
- B. Increase the size of the gp2 volumes in each instance to 3 TB.
- C. Create a new Amazon EFS file system and move all the data to this new file system
- D. Mount this file system to all 10 instances.
- E. Create a new Amazon S3 bucket and move all the data to this new bucket
- F. Allow each instance to access this S3 bucket and use it for storage.

Answer: B

NEW QUESTION 2

A company receives clickstream data files to Amazon S3 every five minutes. A Python script runs as a cron job once a day on an Amazon EC2 instance to process each file and load it into a database hosted on Amazon RDS. The cron job takes 15 to 30 minutes to process 24 hours of data. The data consumers ask for the data to be available as soon as possible.

Which solution would accomplish the desired outcome?

- A. Increase the size of the instance to speed up processing and update the schedule to run once an hour.
- B. Convert the cron job to an AWS Lambda function and trigger this new function using a cron job on an EC2 instance.
- C. Convert the cron job to an AWS Lambda function and schedule it to run once an hour using Amazon CloudWatch events.
- D. Create an AWS Lambda function that runs when a file is delivered to Amazon S3 using S3 event notifications.

Answer: D

Explanation:

<https://docs.aws.amazon.com/lambda/latest/dg/with-s3.html>

NEW QUESTION 3

An organization has two Amazon EC2 instances:

- The first is running an ordering application and an inventory application.
- The second is running a queuing system.

During certain times of the year, several thousand orders are placed per second. Some orders were lost when the queuing system was down. Also, the organization's inventory application has the incorrect quantity of products because some orders were processed twice.

What should be done to ensure that the applications can handle the increasing number of orders?

- A. Put the ordering and inventory applications into their own AWS Lambda function
- B. Have the ordering application write the messages into an Amazon SQS FIFO queue.
- C. Put the ordering and inventory applications into their own Amazon ECS containers and create an Auto Scaling group for each application
- D. Then, deploy the message queuing server in multiple Availability Zones.
- E. Put the ordering and inventory applications into their own Amazon EC2 instances, and create an Auto Scaling group for each application
- F. Use Amazon SQS standard queues for the incoming orders, and implement idempotency in the inventory application.
- G. Put the ordering and inventory applications into their own Amazon EC2 instance
- H. Write the incoming orders to an Amazon Kinesis data stream. Configure AWS Lambda to poll the stream and update the inventory application.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/standard-queues.html>

NEW QUESTION 4

A large company has many business units. Each business unit has multiple AWS accounts for different purposes. The CIO of the company sees that each business unit has data that would be useful to share with other parts of the company. In total, there are about 10 PB of data that needs to be shared with users in 1,000 AWS accounts. The data is proprietary, so some of it should only be available to users with specific job types. Some of the data is used for throughput of intensive workloads, such as simulations. The number of accounts changes frequently because of new initiatives, acquisitions, and divestitures.

A Solutions Architect has been asked to design a system that will allow for sharing data for use in AWS with all of the employees in the company.

Which approach will allow for secure data sharing in a scalable way?

- A. Store the data in a single Amazon S3 bucket
- B. Create an IAM role for every combination of job type and business unit that allows to appropriate read/write access based on object prefixes in the S3 bucket
- C. The roles should have trust policies that allow the business unit's AWS accounts to assume their role
- D. Use IAM in each business unit's AWS account to prevent them from assuming roles for a different job type
- E. Users get credentials to access the data by using AssumeRole from their business unit's AWS account
- F. Users can then use those credentials with an S3 client.
- G. Store the data in a single Amazon S3 bucket
- H. Write a bucket policy that uses conditions to grant read and write access where appropriate, based on each user's business unit and job type
- I. Determine the business unit with the AWS account accessing the bucket and the job type with a prefix in the IAM user's name
- J. Users can access data by using IAM credentials from their business unit's AWS account with an S3 client.
- K. Store the data in a series of Amazon S3 buckets
- L. Create an application running in Amazon EC2 that is integrated with the company's identity provider (IdP) that authenticates users and allows them to download or upload data through the application
- M. The application uses the business unit and job type information in the IdP to control what users can upload and download through the application
- N. The users can access the data through the application's API.

- O. Store the data in a series of Amazon S3 bucket
- P. Create an AWS STS token vending machine that is integrated with the company's identity provider (IdP). When a user logs in, have the token vending machine attach an IAM policy that assumes the role that limits the user's access and/or upload only the data the user is authorized to access
- Q. Users can get credentials by authenticating to the token vending machine's website or API and then use those credentials with an S3 client.

Answer: B

NEW QUESTION 5

A company has an Amazon EC2 deployment that has the following architecture:

- An application tier that contains 8 m4.xlarge instances
- A Classic Load Balancer
- Amazon S3 as a persistent data store

After one of the EC2 instances fails, users report very slow processing of their requests. A Solutions Architect must recommend design changes to maximize system reliability. The solution must minimize costs.

What should the Solution Architect recommend?

- A. Migrate the existing EC2 instances to a serverless deployment using AWS Lambda functions
- B. Change the Classic Load Balancer to an Application Load Balancer
- C. Replace the application tier with m4.large instances in an Auto Scaling group
- D. Replace the application tier with 4 m4.2xlarge instances

Answer: B

Explanation:

By default, connection draining is enabled for Application Load Balancers but must be enabled for Classic Load Balancers. When Connection Draining is enabled and configured, the process of deregistering an instance from an Elastic Load Balancer gains an additional step. For the duration of the configured timeout, the load balancer will allow existing, in-flight requests made to an instance to complete, but it will not send any new requests to the instance. During this time, the API will report the status of the instance as InService, along with a message stating that "Instance deregistration currently in progress." Once the timeout is reached, any remaining connections will be forcibly closed. <https://docs.aws.amazon.com/autoscaling/ec2/userguide/attach-load-balancer-asg.html>
<https://aws.amazon.com/blogs/aws/elb-connection-draining-remove-instances-from-service-with-care/>

NEW QUESTION 6

A company has a standard three-tier architecture using two Availability Zones. During the company's off season, users report that the website is not working. The Solutions Architect finds that no changes have been made to the environment recently, the website is reachable, and it is possible to log in. However, when the Solutions Architect selects the "find a store near you" function, the maps provided on the site by a third-party RESTful API call do not work about 50% of the time after refreshing the page. The outbound API calls are made through Amazon EC2 NAT instances.

What is the MOST likely reason for this failure and how can it be mitigated in the future?

- A. The network ACL for one subnet is blocking outbound web traffic
- B. Open the network ACL and prevent administration from making future changes through IAM.
- C. The fault is in the third-party environment
- D. Contact the third party that provides the maps and request a fix that will provide better uptime.
- E. One NAT instance has become overloaded
- F. Replace both EC2 NAT instances with a larger-sized instance and make sure to account for growth when making the new instance size.
- G. One of the NAT instances failed
- H. Recommend replacing the EC2 NAT instances with a NAT gateway.

Answer: D

Explanation:

The issue is 50% failure, means the balancing over 2 AZs is failing on one NAT instance in one AZ. The solution is to replace the NAT instance with fully managed and high available NAT gateway.

NEW QUESTION 7

A company is designing a new highly available web application on AWS. The application requires consistent and reliable connectivity from the application servers in AWS to a backend REST API hosted in the company's on-premises environment. The backend connection between AWS and on-premises will be routed over an AWS Direct Connect connection through a private virtual interface. Amazon Route 53 will be used to manage private DNS records for the application to resolve the IP address on the backend REST API.

Which design would provide a reliable connection to the backend API?

- A. Implement at least two backend endpoints for the backend REST API, and use Route 53 health checks to monitor the availability of each backend endpoint and perform DNS-level failover.
- B. Install a second Direct Connect connection from a different network carrier and attach it to the same virtual private gateway as the first Direct Connect connection.
- C. Install a second cross connect for the same Direct Connect connection from the same network carrier, and join both connections to the same link aggregation group (LAG) on the same private virtual interface.
- D. Create an IPsec VPN connection routed over the public internet from the on-premises data center to AWS and attach it to the same virtual private gateway as the Direct Connect connection.

Answer: A

NEW QUESTION 8

A company runs a memory-intensive analytics application using on-demand Amazon EC2 compute optimized instance. The application is used continuously and application demand doubles during working hours. The application currently scales based on CPU usage. When scaling in occurs, a lifecycle hook is used because the instance requires 4 minutes to clean the application state before terminating.

Because users reported poor performance during working hours, scheduled scaling actions were implemented so additional instances would be added during working hours. The Solutions Architect has been asked to reduce the cost of the application.

Which solution is MOST cost-effective?

- A. Use the existing launch configuration that uses C5 instances, and update the application AMI to include the Amazon CloudWatch agent
- B. Change the Auto Scaling policies to scale based on memory utilization
- C. Use Reserved Instances for the number of instances required after working hours, and use Spot Instances to cover the increased demand during working hours.
- D. Update the existing launch configuration to use R5 instances, and update the application AMI to include SSM Agent
- E. Change the Auto Scaling policies to scale based on memory utilization
- F. Use Reserved instances for the number of instances required after working hours, and use Spot Instances with On-Demand instances to cover the increased demand during working hours.
- G. Use the existing launch configuration that uses C5 instances, and update the application AMI to include SSM Agent
- H. Leave the Auto Scaling policies to scale based on CPU utilization
- I. Use scheduled Reserved Instances for the number of instances required after working hours, and use Spot Instances to cover the increased demand during work hours.
- J. Create a new launch configuration using R5 instances, and update the application AMI to include the Amazon CloudWatch agent
- K. Change the Auto Scaling policies to scale based on memory utilization
- L. Use Reserved Instances for the number of instances required after working hours, and use Standard Reserved Instances with On-Demand Instances to cover the increased demand during working hours.

Answer: D

Explanation:

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/monitoring_ec2.html

NEW QUESTION 9

While debugging a backend application for an IoT system that supports globally distributed devices a Solutions Architect notices that stale data is occasionally being sent to user devices. Devices often share data, and stale data does not cause issues in most cases However device operations are disrupted when a device reads the stale data after an update

The global system has multiple identical application stacks deployed in different AWS Regions If a user device travels out of its home geographic region it will always connect to the geographically closest AWS Region to write or read data The same data is available in all supported AWS Regions using an Amazon DynamoDB global table

What change should be made to avoid causing disruptions in device operations'?

- A. Update the backend to use strongly consistent read
- B. Update the devices to always write to and read from their home AWS Region
- C. Enable strong consistency globally on a DynamoDB global table Update the backend to use strongly consistent reads
- D. Switch the backend data store to Amazon Aurora MySQL with cross-region replicas Update the backend to always write to the master endpoint
- E. Select one AWS Region as a master and perform all writes in that AWS Region only Update the backend to use strongly consistent reads

Answer: B

NEW QUESTION 10

A company has a serverless application comprised of Amazon CloudFront, Amazon API Gateway, and AWS Lambda functions. The current deployment process of the application code is to create a new version number of the Lambda function and run an AWS CLI script to update. If the new function version has errors, another CLI script reverts by deploying the previous working version of the function. The company would like to decrease the time to deploy new versions of the application logic provided by the Lambda functions, and also reduce the time to detect and revert when errors are identified.

How can this be accomplished?

- A. Create and deploy nested AWS CloudFormation stacks with the parent stack consisting of the AWS CloudFront distribution and API Gateway, and the child stack containing the Lambda function
- B. For changes to Lambda, create an AWS CloudFormation change set and deploy; if errors are triggered, revert the AWS CloudFormation change set to the previous version.
- C. Use AWS SAM and built-in AWS CodeDeploy to deploy the new Lambda version, gradually shift traffic to the new version, and use pre-traffic and post-traffic test functions to verify code
- D. Rollback if Amazon CloudWatch alarms are triggered.
- E. Refactor the AWS CLI scripts into a single script that deploys the new Lambda version
- F. When deployment is completed, the script tests execution
- G. If errors are detected, revert to the previous Lambda version.
- H. Create and deploy an AWS CloudFormation stack that consists of a new API Gateway endpoint that references the new Lambda version
- I. Change the CloudFront origin to the new API Gateway endpoint, monitor errors and if detected, change the AWS CloudFront origin to the previous API Gateway endpoint.

Answer: B

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2017/11/aws-lambda-supports-traffic-shifting-and-phased-deploy> <https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/automating-updates-to-serverless>

NEW QUESTION 10

A company runs an application on a fleet of Amazon EC2 instances The application requires low latency and random access to 100 GB of data The application must be able to access the data at up to 3,000 IOPS A Development team has configured the EC2 launch template to provision a 100-GB Provisioned IOPS (PIOPS) Amazon EBS volume with 3,000 IOPS provisioned A Solutions Architect is tasked with lowering costs without impacting performance and durability Which action should be taken?

- A. Create an Amazon EFS file system with the performance mode set to Max I/O Configure the EC2 operating system to mount the EFS file system
- B. Create an Amazon EFS file system with the throughput mode set to Provisioned Configure the EC2 operating system to mount the EFS file system
- C. Update the EC2 launch template to allocate a new 1-TB EBS General Purpose SSO (gp2) volume
- D. Update the EC2 launch template to exclude the PIOPS volume Configure the application to use local instance storage

Answer: A

NEW QUESTION 14

As a part of building large applications in the AWS Cloud, the Solutions Architect is required to implement the perimeter security protection. Applications running on AWS have the following endpoints:

- Application Load Balancer
- Amazon API Gateway regional endpoint
- Elastic IP address-based EC2 instances.
- Amazon S3 hosted websites.
- Classic Load Balancer

The Solutions Architect must design a solution to protect all of the listed web front ends and provide the following security capabilities:

- DDoS protection
- SQL injection protection
- IP address whitelist/blacklist
- HTTP flood protection
- Bad bot scraper protection

How should the Solutions Architect design the solution?

- A. Deploy AWS WAF and AWS Shield Advanced on all web endpoint
- B. Add AWS WAF rules to enforce the company's requirements.
- C. Deploy Amazon CloudFront in front of all the endpoint
- D. The CloudFront distribution provides perimeter protection
- E. Add AWS Lambda-based automation to provide additional security.
- F. Deploy Amazon CloudFront in front of all the endpoint
- G. Deploy AWS WAF and AWS Shield Advance
- H. Add AWS WAF rules to enforce the company's requirement
- I. Use AWS Lambda to automate and enhance the security posture.
- J. Secure the endpoints by using network ACLs and security groups and adding rules to enforce the company's requirement
- K. Use AWS Lambda to automatically update the rules.

Answer: C

NEW QUESTION 16

A company CFO recently analyzed the company's AWS monthly bill and identified an opportunity to reduce the cost for AWS Elastic Beanstalk environments in use. The CFO has asked a Solutions Architect to design a highly available solution that will spin up an Elastic Beanstalk environment in the morning and terminate it at the end of the day.

The solution should be designed with minimal operational overhead and to minimize costs. It should also be able to handle the increased use of Elastic Beanstalk environments among different teams, and must provide a one-stop scheduler solution for all teams to keep the operational costs low.

What design will meet these requirements?

- A. Set up a Linux EC2 Micro instance
- B. Configure an IAM role to allow the start and stop of the Elastic Beanstalk environment and attach it to the instance
- C. Create scripts on the instance to start and stop the Elastic Beanstalk environment
- D. Configure cron jobs on the instance to execute the scripts.
- E. Develop AWS Lambda functions to start and stop the Elastic Beanstalk environment
- F. Configure a Lambda execution role granting Elastic Beanstalk environment start/stop permissions, and assign the role to the Lambda function
- G. Configure cron expression Amazon CloudWatch Events rules to trigger the Lambda functions.
- H. Develop an AWS Step Functions state machine with "wait" as its type to control the start and stop time. Use the activity task to start and stop the Elastic Beanstalk environment
- I. Create a role for Step Functions to allow it to start and stop the Elastic Beanstalk environment
- J. Invoke Step Functions daily.
- K. Configure a time-based Auto Scaling group
- L. In the morning, have the Auto Scaling group scale up an Amazon EC2 instance and put the Elastic Beanstalk environment start command in the EC2 instance user data
- M. At the end of the day, scale down the instance number to 0 to terminate the EC2 instance.

Answer: B

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/schedule-elastic-beanstalk-stop-restart/>

NEW QUESTION 18

A company is migrating its marketing website and content management system from an on-premises data center to AWS. The company wants the AWS application to be developed in a VPC with Amazon EC2 instances used for the web servers and an Amazon RDS instance for the database.

The company has a runbook document that describes the installation process of the on-premises system. The company would like to base the AWS system on the processes referenced in the runbook document. The runbook document describes the installation and configuration of the operating systems, network settings, the website, and content management system software on the servers. After the migration is complete, the company wants to be able to make changes quickly to take advantage of other AWS features.

How can the application and environment be deployed and automated in AWS, while allowing for future changes?

- A. Update the runbook to describe how to create the VPC, the EC2 instances, and the RDS instance for the application by using the AWS Console
- B. Make sure that the rest of the steps in the runbook are updated to reflect any changes that may come from the AWS migration.
- C. Write a Python script that uses the AWS API to create the VPC, the EC2 instances, and the RDS instance for the application
- D. Write shell scripts that implement the rest of the steps in the runbook
- E. Have the Python script copy and run the shell scripts on the newly created instances to complete the installation.
- F. Write an AWS CloudFormation template that creates the VPC, the EC2 instances, and the RDS instance for the application

- G. Ensure that the rest of the steps in the runbook are updated to reflect any changes that may come from the AWS migration.
- H. Write an AWS CloudFormation template that creates the VPC, the EC2 instances, and the RDS instance for the applicatio
- I. Include EC2 user data in the AWS CloudFormation template to install and configure the software.

Answer: D

NEW QUESTION 23

A company has a 24 TB MySQL database in its on-premises data center that grows at the rate of 10 GB per day. The data center is connected to the company's AWS infrastructure with a 50 Mbps VPN connection.

The company is migrating the application and workload to AWS. The application code is already installed and tested on Amazon EC2. The company now needs to migrate the database and wants to go live on AWS within 3 weeks.

Which of the following approaches meets the schedule with LEAST downtime?

- A. 1. Use the VM Import/Export service to import a snapshot on the on-premises database into AWS.2.Launch a new EC2 instance from the snapshot.3. Set up ongoing database replication from on premises to the EC2 database over the VPN.4. Change the DNS entry to point to the EC2 database.5. Stop the replication.
- B. 1. Launch an AWS DMS instance.2. Launch an Amazon RDS Aurora MySQL DB instance.3. Configure the AWS DMS instance with on-premises and Amazon RDS database information.4. Start the replication task within AWS DMS over the VPN.5. Change the DNS entry to point to the Amazon RDS MySQL database.6. Stop the replication.
- C. 1. Create a database export locally using database-native tools.2. Import that into AWS using AWS Snowball.3. Launch an Amazon RDS Aurora DB instance.4. Load the data in the RDS Aurora DB instance from the export.5. Set up database replication from the on-premises database to the RDS Aurora DB instance over the VPN.6. Change the DNS entry to point to the RDS Aurora DB instance.7. Stop the replication.
- D. 1. Take the on-premises application offline.2. Create a database export locally using database-native tools.3. Import that into AWS using AWS Snowball.4. Launch an Amazon RDS Aurora DB instance.5. Load the data in the RDS Aurora DB instance from the export.6. Change the DNS entry to point to the Amazon RDS Aurora DB instance.7. Put the Amazon EC2 hosted application online.

Answer: C

NEW QUESTION 28

A large company has increased its utilization of AWS over time in an unmanaged way. As such, they have a large number of independent AWS accounts across different business units, projects, and environments. The company has created a Cloud Center of Excellence team, which is responsible for managing all aspects of the AWS Cloud, including their AWS accounts.

Which of the following should the Cloud Center of Excellence team do to BEST address their requirements in a centralized way? (Select two.)

- A. Control all AWS account root user credential
- B. Assign AWS IAM users in the account of each user who needs to access AWS resource
- C. Follow the policy of least privilege in assigning permissions to each user.
- D. Tag all AWS resources with details about the business unit, project, and environmen
- E. Send all AWS Cost and Usage reports to a central Amazon S3 bucket, and use tools such as Amazon Athena and Amazon QuickSight to collect billing details by business unit.
- F. Use the AWS Marketplace to choose and deploy a Cost Management too
- G. Tag all AWS resources with details about the business unit, project, and environmen
- H. Send all AWS Cost and Usage reports for the AWS accounts to this tool for analysis.
- I. Set up AWS Organization
- J. Enable consolidated billing, and link all existing AWS accounts to a master billing account
- K. Tag all AWS resources with details about the business unit, project and environmen
- L. Analyze Cost and Usage reports using tools such as Amazon Athena and Amazon QuickSight to collect billing details by business unit.
- M. Using a master AWS account, create IAM users within the master account
- N. Define IAM roles in the other AWS accounts, which cover each of the required functions in the account
- O. Follow the policy of least privilege in assigning permissions to each role, then enable the IAM users to assume the roles that they need to use.

Answer: DE

NEW QUESTION 31

A company is using AWS CloudFormation to deploy its infrastructure. The company is concerned that, if a production CloudFormation stack is deleted, important data stored in Amazon RDS databases or Amazon EBS volumes might also be deleted.

How can the company prevent users from accidentally deleting data in this way?

- A. Modify the CloudFormation templates to add a DeletionPolicy attribute to RDS and EBS resources.
- B. Configure a stack policy that disallows the deletion of RDS and EBS resources.
- C. Modify IAM policies to deny deleting RDS and EBS resources that are tagged with an "aws:cloudformation:stack-name" tag.
- D. Use AWS Config rules to prevent deleting RDS and EBS resources.

Answer: A

Explanation:

With the DeletionPolicy attribute you can preserve or (in some cases) backup a resource when its stack is deleted. You specify a DeletionPolicy attribute for each resource that you want to control. If a resource has no DeletionPolicy attribute, AWS CloudFormation deletes the resource by default. To keep a resource when its stack is deleted, specify Retain for that resource. You can use retain for any resource. For example, you can retain a nested stack, Amazon S3 bucket, or EC2 instance so that you can continue to use or modify those resources after you delete their stacks.

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-deletionpolicy.html>

NEW QUESTION 36

A company will 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": "2012-10-27",
  "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

NEW QUESTION 40

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

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

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

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

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

Answer: B

NEW QUESTION 42

A company uses Amazon S3 to store documents that may only be accessible to an Amazon EC2 instance in a certain virtual private cloud (VPC). The company fears that a malicious insider with access to this instance could also set up an EC2 instance in another VPC to access these documents.

Which of the following solutions will provide the required protection?

- A. Use an S3 VPC endpoint and an S3 bucket policy to limit access to this VPC endpoint.
- B. Use EC2 instance profiles and an S3 bucket policy to limit access to the role attached to the instance profile.
- C. Use S3 client-side encryption and store the key in the instance metadata.
- D. Use S3 server-side encryption and protect the key with an encryption context.

Answer: A

Explanation:

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

Endpoint connections cannot be extended out of a VPC. Resources on the other side of a VPN connection, VPC peering connection, AWS Direct Connect connection, or ClassicLink connection in your VPC cannot use the endpoint to communicate with resources in the endpoint service.

NEW QUESTION 46

A company is migrating a subset of its application APIs from Amazon EC2 instances to run on a serverless infrastructure. The company has set up Amazon API Gateway, AWS Lambda, and Amazon DynamoDB for the new application. The primary responsibility of the Lambda function is to obtain data from a third-party Software as a Service (SaaS) provider. For consistency, the Lambda function is attached to the same virtual private cloud (VPC) as the original EC2 instances. Test users report an inability to use this newly moved functionality, and the company is receiving 5xx errors from API Gateway. Monitoring reports from the SaaS provider shows that the requests never made it to its systems. The company notices that Amazon CloudWatch Logs are being generated by the Lambda functions. When the same functionality is tested against the EC2 systems, it works as expected. What is causing the issue?

- A. Lambda is in a subnet that does not have a NAT gateway attached to it to connect to the SaaS provider.
- B. The end-user application is misconfigured to continue using the endpoint backed by EC2 instances.
- C. The throttle limit set on API Gateway is too low and the requests are not making their way through.
- D. API Gateway does not have the necessary permissions to invoke Lambda.

Answer: A

NEW QUESTION 49

A company has an existing on-premises three-tier web application. The Linux web servers serve content from a centralized file share on a NAS server because the content is refreshed several times a day from various sources. The existing infrastructure is not optimized and the company would like to move to AWS in order to gain the ability to scale resources up and down in response to load. On-premises and AWS resources are connected using AWS Direct Connect. How can the company migrate the web infrastructure to AWS without delaying the content refresh process?

- A. Create a cluster of web server Amazon EC2 instances behind a Classic Load Balancer on AW
- B. Share an Amazon EBS volume among all instances for the conten
- C. Schedule a periodic synchronization of this volume and the NAS server.
- D. Create an on-premises file gateway using AWS Storage Gateway to replace the NAS server and replicate content to AW
- E. On the AWS side, mount the same Storage Gateway bucket to each web server Amazon EC2 instance to serve the content.
- F. Expose an Amazon EFS share to on-premises users to serve as the NAS serv
- G. Mount the same EFS share to the web server Amazon EC2 instances to serve the content.
- H. Create web server Amazon EC2 instances on AWS in an Auto Scaling grou
- I. Configure a nightly process where the web server instances are updated from the NAS server.

Answer: C

Explanation:

File gateway is limited by performance its gateway instance, whether EC2 or On-premises, Cache will get filled up fast if not properly configured, For large number of EC2 instances EFS scales better. So, bottom line is File Storage gateway is for legacy applications and you have to add cost of large gateway instances before comparing it to same quantity of EFS storage. https://www.reddit.com/r/aws/comments/82pyop/storage_gateway_vs_efs/
<https://docs.aws.amazon.com/efs/latest/ug/efs-onpremises.html>

NEW QUESTION 53

A development team has created a series of AWS CloudFormation templates to help deploy services. They created a template for a network/virtual private (VPC) stack, a database stack, a bastion host stack, and a web application-specific stack. Each service requires the deployment of at least: Each template has multiple input parameters that make it difficult to deploy the services individually from the AWS CloudFormation console. The input parameters from one stack are typically outputs from other stacks. For example, the VPC ID, subnet IDs, and security groups from the network stack may need to be used in the application stack or database stack.

Which actions will help reduce the operational burden and the number of parameters passed into a service deployment? (Choose two.)

- A. Create a new AWS CloudFormation template for each servic
- B. After the existing templates to use cross-stack references to eliminate passing many parameters to each templat
- C. Call each required stack for the application as a nested stack from the new stac
- D. Call the newly created service stack from theAWS CloudFormation console to deploy the specific service with a subset of the parameters previously required.
- E. Create a new portfolio in AWS Service Catalog for each servic
- F. Create a product for each existing AWS CloudFormation template required to build the servic
- G. Add the products to the portfolio that represents that service in AWS Service Catalo
- H. To deploy the service, select the specific service portfolio and launch the portfolio with the necessary parameters to deploy all templates.
- I. Set up an AWS CodePipeline workflow for each servic
- J. For each existing template, choose AWS CloudFormation as a deployment actio
- K. Add the AWS CloudFormation template to the deployment actio
- L. Ensure that the deployment actions are processed to make sure that dependences are obeye
- M. Use configuration files and scripts to share parameters between the stack
- N. To launch the service, execute the specific template by choosing the name of the service and releasing a change.
- O. Use AWS Step Functions to define a new servic
- P. Create a new AWS CloudFormation template for each servic
- Q. After the existing templates to use cross-stack references to eliminate passing many parameters to each templat
- R. Call each required stack for the application as a nested stack from the new service templat
- S. Configure AWS Step Functions to call the service template directl
- T. In the AWS Step Functions console, execute the step.
- . Create a new portfolio for the Services in AWS Service Catalo
- . Create a new AWS CloudFormation template for each servic
- . After the existing templates to use cross-stack references to eliminate passing many parameters to each templat
- . Call each required stack for the application as a nested stack from the new stac
- . Create a product for each applicatio
- . Add the service template to the produc
- . Add each new product to the portfoli
- . Deploy the product from the portfolio to deploy the service with the necessary parameters only to start the deployment.

Answer: AE

NEW QUESTION 55

A bank is designing an online customer service portal where customers can chat with customer service agents. The portal is required to maintain a 15-minute RPO or RTO in case of a regional disaster. Banking regulations require that all customer service chat transcripts must be preserved on durable storage for at least 7 years, chat conversations must be encrypted in-flight, and transcripts must be encrypted at rest. The Data Lost Prevention team requires that data at rest must be encrypted using a key that the team controls, rotates, and revokes.

Which design meets these requirements?

- A. The chat application logs each chat message into Amazon CloudWatch Log
- B. A scheduled AWS Lambda function invokes a CloudWatch Log
- C. CreateExportTask every 5 minutes to export chat transcripts to Amazon S3. The S3 bucket is configured for cross-region replication to the backup regio
- D. Separate AWS KMS keys are specified for the CloudWatch Logs group and the S3 bucket.
- E. The chat application logs each chat message into two different Amazon CloudWatch Logs groups in two different regions, with the same AWS KMS key applie
- F. Both CloudWatch Logs groups are configured to export logs into an Amazon Glacier vault with a 7-year vault lock policy with a KMS key specified.
- G. The chat application logs each chat message into Amazon CloudWatch Log
- H. A subscription filter on the CloudWatch Logs group feeds into an Amazon Kinesis Data Firehose which streams the chat messages into an Amazon S3 bucket in the backup regio
- I. Separate AWS KMS keys are specified for the CloudWatch Logs group and the Kinesis Data Firehose.
- J. The chat application logs each chat message into Amazon CloudWatch Log
- K. The CloudWatch Logs group is configured to export logs into an Amazon Glacier vault with a 7-year vault lock polic
- L. Glacier cross-region replication mirrors chat archives to the backup regio
- M. Separate AWS KMS keys are specified for the CloudWatch Logs group and the Amazon Glacier vault.

Answer: B

NEW QUESTION 59

A Development team is deploying new APIs as serverless applications within a company. The team is currently using the AWS Management Console to provision Amazon API Gateway, AWS Lambda, and Amazon DynamoDB resources. A Solutions Architect has been tasked with automating the future deployments of these serverless APIs.

How can this be accomplished?

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

Answer: B

Explanation:

<https://aws-quickstart.s3.amazonaws.com/quickstart-trek10-serverless-enterprise-cicd/doc/serverless-cicd-for-th>
<https://aws.amazon.com/quickstart/architecture/serverless-cicd-for-enterprise/>

NEW QUESTION 60

A three-tier web application runs on Amazon EC2 instances. Cron daemons are used to trigger scripts that collect the web server, application, and database logs and send them to a centralized location every hour. Occasionally, scaling events or unplanned outages have caused the instances to stop before the latest logs were collected, and the log files were lost.

Which of the following options is the MOST reliable way of collecting and preserving the log files?

- A. Update the cron jobs to run every 5 minutes instead of every hour to reduce the possibility of log messages being lost in an outage.
- B. Use Amazon CloudWatch Events to trigger Amazon Systems Manager Run Command to invoke the log collection scripts more frequently to reduce the possibility of log messages being lost in an outage.
- C. Use the Amazon CloudWatch Logs agent to stream log messages directly to CloudWatch Logs. Configure the agent with a batch count of 1 to reduce the possibility of log messages being lost in an outage.
- D. Use Amazon CloudWatch Events to trigger AWS Lambda to SSH into each running instance and invoke the log collection scripts more frequently to reduce the possibility of log messages being lost in an outage.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/AgentReference.html>

NEW QUESTION 65

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

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

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

- A. An inbound rule for port 80 from source 0.0.0 0/0
- B. An inbound rule for port 80 from source 10.0.0.0/24

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

Answer: BE

NEW QUESTION 67

A large multinational company runs a timesheet application on AWS that is used by staff across the world. The application runs on Amazon EC2 instances in an Auto Scaling group behind an Elastic Load Balancing (ELB) load balancer, and stores in an Amazon RDS MySQL Multi-AZ database instance. The CFO is concerned about the impact on the business if the application is not available. The application must not be down for more than two hours, but the solution must be as cost-effective as possible.

How should the Solutions Architect meet the CFO's requirements while minimizing data loss?

- A. In another region, configure a read replica and create a copy of the infrastructure
- B. When an issue occurs, promote the read replica and configure as an Amazon RDS Multi-AZ database instance
- C. Update the DNS to point to the other region's ELB.
- D. Configure a 1-day window of 60-minute snapshots of the Amazon RDS Multi-AZ database instance. Create an AWS CloudFormation template of the application infrastructure that uses the latest snapshot
- E. When an issue occurs, use the AWS CloudFormation template to create the environment in another region
- F. Update the DNS record to point to the other region's ELB.
- G. Configure a 1-day window of 60-minute snapshots of the Amazon RDS Multi-AZ database instance which is copied to another region
- H. Create an AWS CloudFormation template of the application infrastructure that uses the latest copied snapshot
- I. When an issue occurs, use the AWS CloudFormation template to create the environment in another region
- J. Update the DNS record to point to the other region's ELB.
- K. Configure a read replica in another region
- L. Create an AWS CloudFormation template of the application infrastructure
- M. When an issue occurs, promote the read replica and configure as an Amazon RDS Multi-AZ database instance and use the AWS CloudFormation template to create the environment in another region using the promoted Amazon RDS instance
- N. Update the DNS record to point to the other region's ELB.

Answer: D

NEW QUESTION 69

A company is using AWS for production and development workloads. Each business unit has its own AWS account for production, and a separate AWS account to develop and deploy its applications. The Information Security department has introduced new security policies that limit access for terminating certain Amazon EC2 instances in all accounts to a small group of individuals from the Security team.

How can the Solutions Architect meet these requirements?

- A. Create a new IAM policy that allows access to those EC2 instances only for the Security team
- B. Apply this policy to the AWS Organizations master account.
- C. Create a new tag-based IAM policy that allows access to these EC2 instances only for the Security team. Tag the instances appropriately, and apply this policy in each account.
- D. Create an organizational unit under AWS Organization
- E. Move all the accounts into this organizational unit and use SCP to apply a whitelist policy to allow access to these EC2 instances for the Security team only.
- F. Set up SAML federation for all accounts in AWS
- G. Configure SAML so that it checks for the service API call before authenticating the user
- H. Block SAML from authenticating API calls if anyone other than the Security team accesses these instances.

Answer: B

NEW QUESTION 70

A company has an application that uses Amazon EC2 instances in an Auto Scaling group. The Quality Assurance (QA) department needs to launch a large number of short-lived environments to test the application. The application environments are currently launched by the Manager of the department using an AWS CloudFormation template. To launch the stack, the Manager uses a role with permission to use CloudFormation, EC2 and Auto Scaling APIs. The Manager wants to allow testers to launch their own environments, but does not want to grant broad permission to each user. Which set up would achieve these goals?

- A. Upload the AWS CloudFormation template to Amazon S3. Give users in the QA department permission to assume the Manager's role and add a policy that restricts the permissions to the template and the resources it creates
- B. Train users to launch the template from the CloudFormation console.
- C. Create an AWS Service Catalog product from the environment template
- D. Add a launch constraint to the product with the existing role
- E. Give users in the QA department permission to use AWS Service Catalog APIs only
- F. Train users to launch the templates from the AWS Service Catalog console.
- G. Upload the AWS CloudFormation template to Amazon S3. Give users in the QA department permission to use CloudFormation and S3 APIs, with conditions that restrict the permission to the template and the resources it creates
- H. Train users to launch the template from the CloudFormation console.
- I. Create an AWS Elastic Beanstalk application from the environment template
- J. Give users in the QA department permission to use Elastic Beanstalk permissions only
- K. Train users to launch Elastic Beanstalk environments with the Elastic Beanstalk CLI, passing the existing role to the environment as a service role.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/mt/how-to-launch-secure-and-governed-aws-resources-with-aws-cloudformation->

NEW QUESTION 75

A Solutions Architect is redesigning an image-viewing and messaging platform to be delivered as SaaS. Currently, there is a farm of virtual desktop infrastructure (VDI) that runs a desktop image-viewing application and a desktop messaging application. Both applications use a shared database to manage user accounts and sharing. Users log in from a web portal that launches the applications and streams the view of the application on the user's machine. The Development

Operations team wants to move away from using VDI and wants to rewrite the application.
 What is the MOST cost-effective architecture that offers both security and ease of management?

- A. Run a website from an Amazon S3 bucket with a separate S3 bucket for images and messaging data. Call AWS Lambda functions from embedded JavaScript to manage the dynamic content, and use Amazon Cognito for user and sharing management.
- B. Run a website from Amazon EC2 Linux servers, storing the images in Amazon S3, and use Amazon Cognito for user accounts and sharing.
- C. Create AWS CloudFormation templates to launch the application by using EC2 user data to install and configure the application.
- D. Run a website as an AWS Elastic Beanstalk application, storing the images in Amazon S3, and using an Amazon RDS database for user accounts and sharing.
- E. Create AWS CloudFormation templates to launch the application and perform blue/green deployments.
- F. Run a website from an Amazon S3 bucket that authorizes Amazon AppStream to stream applications for a combined image viewer and messenger that stores images in Amazon S3. Have the website use an Amazon RDS database for user accounts and sharing.

Answer: D

Explanation:

<https://docs.aws.amazon.com/appstream2/latest/developerguide/managing-images.html>

NEW QUESTION 77

A Solutions Architect is designing the storage layer for a recently purchased application. The application will be running on Amazon EC2 instances and has the following layers and requirements:

- Data layer: A POSIX file system shared across many systems.
- Service layer: Static file content that requires block storage with more than 100k IOPS. Which combination of AWS services will meet these needs? (Choose two.)

- A. Data layer – Amazon S3
- B. Data layer – Amazon EC2 Ephemeral Storage
- C. Data layer – Amazon EFS
- D. Service layer – Amazon EBS volumes with Provisioned IOPS
- E. Service layer – Amazon EC2 Ephemeral Storage

Answer: CE

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/storage-optimized-instances.html>

NEW QUESTION 81

A Solutions Architect must update an application environment within AWS Elastic Beanstalk using a blue/green deployment methodology. The Solutions Architect creates an environment that is identical to the existing application environment and deploys the application to the new environment. What should be done next to complete the update?

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

Answer: B

Explanation:

<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.CNAMEswap.html>

NEW QUESTION 86

A company runs an ordering system on AWS using Amazon SQS and AWS Lambda, with each order received as a JSON message. Recently the company had a marketing event that led to a tenfold increase in orders. With this increase, the following undesired behaviors started in the ordering system:

- Lambda failures while processing orders lead to queue backlogs.
- The same orders have been processed multiple times.

A solutions Architect has been asked to solve the existing issues with the ordering system and add the following resiliency features:

- Retain problematic orders for analysis.
- Send notification if errors go beyond a threshold value. How should the Solutions Architect meet these requirements?

- A. Receive multiple messages with each Lambda invocation, add error handling to message processing code and delete messages after processing, increase the visibility timeout for the messages, create a dead letter queue for messages that could not be processed, create an Amazon CloudWatch alarm on Lambda errors for notification.
- B. Receive single messages with each Lambda invocation, put additional Lambda workers to poll the queue, delete messages after processing, increase the message timer for the messages, use Amazon CloudWatch Logs for messages that could not be processed, create a CloudWatch alarm on Lambda errors for notification.
- C. Receive multiple messages with each Lambda invocation, use long polling when receiving the messages, log the errors from the message processing code using Amazon CloudWatch Logs, create a dead letter queue with AWS Lambda to capture failed invocations, create CloudWatch events on Lambda errors for notification.
- D. Receive multiple messages with each Lambda invocation, add error handling to message processing code and delete messages after processing, increase the visibility timeout for the messages, create a delay queue for messages that could not be processed, create an Amazon CloudWatch metric on Lambda errors for notification.

Answer: D

NEW QUESTION 87

A Solutions Architect has created an AWS CloudFormation template for a three-tier application that contains an Auto Scaling group of Amazon EC2 instances running a custom AMI.
 The Solutions Architect wants to ensure that future updates to the custom AMI can be deployed to a running stack by first updating the template to refer to the new AMI, and then invoking UpdateStack to replace the EC2 instances with instances launched from the new AMI.
 How can updates to the AMI be deployed to meet these requirements?

- A. Create a change set for a new version of the template, view the changes to the running EC2 instances to ensure that the AMI is correctly updated, and then execute the change set.
- B. Edit the AWS::AutoScaling::LaunchConfiguration resource in the template, changing its DeletionPolicy to Replace.
- C. Edit the AWS::AutoScaling::AutoScalingGroup resource in the template, inserting an UpdatePolicy attribute.
- D. Create a new stack from the updated template.
- E. Once it is successfully deployed, modify the DNS records to point to the new stack and delete the old stack.

Answer: C

Explanation:

References:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-properties-as-launchconfig.html>

NEW QUESTION 90

A company runs a legacy system on a single m4.2xlarge Amazon EC2 instance with Amazon EBS2 storage. The EC2 instance runs both the web server and a self-managed Oracle database. A snapshot is made of the EBS volume every 12 hours, and an AMI was created from the fully configured EC2 instance. A recent event that terminated the EC2 instance led to several hours of downtime. The application was successfully launched from the AMI, but the age of the EBS snapshot and the repair of the database resulted in the loss of 8 hours of data. The system was also down for 4 hours while the Systems Operators manually performed these processes.
 What architectural changes will minimize downtime and reduce the chance of lost data?

- A. Create an Amazon CloudWatch alarm to automatically recover the instance.
- B. Create a script that will check and repair the database upon reboot.
- C. Subscribe the Operations team to the Amazon SNS message generated by the CloudWatch alarm.
- D. Run the application on m4.xlarge EC2 instances behind an Elastic Load Balancer/Application Load Balance.
- E. Run the EC2 instances in an Auto Scaling group across multiple Availability Zones with a minimum instance count of two.
- F. Migrate the database to an Amazon RDS Oracle Multi-AZ DB instance.
- G. Run the application on m4.2xlarge EC2 instances behind an Elastic Load Balancer/Application Load Balance.
- H. Run the EC2 instances in an Auto Scaling group across multiple Availability Zones with a minimum instance count of one.
- I. Migrate the database to an Amazon RDS Oracle Multi-AZ DB instance.
- J. Increase the web server instance count to two m4.xlarge instances and use Amazon Route 53 round-robin load balancing to spread the load.
- K. Enable Route 53 health checks on the web server.
- L. Migrate the database to an Amazon RDS Oracle Multi-AZ DB instance.

Answer: B

Explanation:

Ensures that there are at least two EC instances, each of which is in a different AZ. It also ensures that the database spans multiple AZs. Hence this meets all the criteria.

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html>

NEW QUESTION 94

What combination of steps could a Solutions Architect take to protect a web workload running on Amazon EC2 from DDoS and application layer attacks? (Select two.)

- A. Put the EC2 instances behind a Network Load Balancer and configure AWS WAF on it.
- B. Migrate the DNS to Amazon Route 53 and use AWS Shield.
- C. Put the EC2 instances in an Auto Scaling group and configure AWS WAF on it.
- D. Create and use an Amazon CloudFront distribution and configure AWS WAF on it.
- E. Create and use an internet gateway in the VPC and use AWS Shield.

Answer: BD

Explanation:

References: <https://aws.amazon.com/answers/networking/aws-ddos-attack-mitigation/>

NEW QUESTION 98

A media company has a 30-TB repository of digital news videos. These videos are stored on tape in an on-premises tape library and referenced by a Media Asset Management (MAM) system. The company wants to enrich the metadata for these videos in an automated fashion and put them into a searchable catalog by using a MAM feature. The company must be able to search based on information in the video, such as objects, scenery items, or people's faces. A catalog is available that contains faces of people who have appeared in the videos that include an image of each person. The company would like to migrate these videos to AWS.
 The company has a high-speed AWS Direct Connect connection with AWS and would like to move the MAM solution video content directly from its current file system.
 How can these requirements be met by using the LEAST amount of ongoing management overhead and causing MINIMAL disruption to the existing system?

- A. Set up an AWS Storage Gateway, file gateway appliance on-premise.
- B. Use the MAM solution to extract the videos from the current archive and push them into the file gateway.
- C. Use the catalog of faces to build a collection in Amazon Rekognition.
- D. Build an AWS Lambda function that invokes the Rekognition Javascript SDK to have Rekognition pull the video from the Amazon S3 files backing the file gateway, retrieve the required metadata, and push the metadata into the MAM solution.
- E. Set up an AWS Storage Gateway, tape gateway appliance on-premise.
- F. Use the MAM solution to extract the videos from the current archive and push them into the tape gateway.
- G. Use the catalog of faces to build a collection in Amazon Rekognition.

- H. Build an AWS Lambda function that invokes the Rekognition Javascript SDK to have Amazon Rekognition process the video in the tape gateway, retrieve the required metadata, and push the metadata into the MAM solution.
- I. Configure a video ingestion stream by using Amazon Kinesis Video Stream
- J. Use the catalog of faces to build a collection in Amazon Rekognition
- K. Stream the videos from the MAM solution into Kinesis Video Stream
- L. Configure Amazon Rekognition to process the streamed video
- M. Then, use a stream consumer to retrieve the required metadata, and push the metadata into the MAM solution
- N. Configure the stream to store the videos in Amazon S3.
- O. Set up an Amazon EC2 instance that runs the OpenCV libraries
- P. Copy the videos, images, and face catalog from the on-premises library into an Amazon EBS volume mounted on this EC2 instance
- Q. Process the videos to retrieve the required metadata, and push the metadata into the MAM solution while also copying the video files to an Amazon S3 bucket.

Answer: C

Explanation:

<https://docs.aws.amazon.com/rekognition/latest/dg/streaming-video.html>

NEW QUESTION 101

A group of research institutions and hospitals are in a partnership to study 2 PBs of genomic data. The institute that owns the data stores it in an Amazon S3 bucket and updates it regularly. The institute would like to give all of the organizations in the partnership read access to the data. All members of the partnership are extremely cost-conscious, and the institute that owns the account with the S3 bucket is concerned about covering the costs for requests and data transfers from Amazon S3.

Which solution allows for secure datasharing without causing the institute that owns the bucket to assume all the costs for S3 requests and data transfers?

- A. Ensure that all organizations in the partnership have AWS account
- B. In the account with the S3 bucket, create a cross-account role for each account in the partnership that allows read access to the data
- C. Have the organizations assume and use that read role when accessing the data.
- D. Ensure that all organizations in the partnership have AWS account
- E. Create a bucket policy on the bucket that owns the data
- F. The policy should allow the accounts in the partnership read access to the bucket
- G. Enable Requester Pays on the bucket
- H. Have the organizations use their AWS credentials when accessing the data.
- I. Ensure that all organizations in the partnership have AWS account
- J. Configure buckets in each of the accounts with a bucket policy that allows the institute that owns the data the ability to write to the bucket
- K. Periodically sync the data from the institute's account to the other organization
- L. Have the organizations use their AWS credentials when accessing the data using their accounts.
- M. Ensure that all organizations in the partnership have AWS account
- N. In the account with the S3 bucket, create a cross-account role for each account in the partnership that allows read access to the data
- O. Enable Requester Pays on the bucket
- P. Have the organizations assume and use that read role when accessing the data.

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/dev/RequesterPaysBuckets.html>

NEW QUESTION 102

A large company experienced a drastic increase in its monthly AWS spend. This is after Developers accidentally launched Amazon EC2 instances in unexpected regions. The company has established practices around least privileges for Developers and controls access to on-premises resources using Active Directory groups. The company now wants to control costs by restricting the level of access that Developers have to the AWS Management Console without impacting their productivity. The company would also like to allow Developers to launch Amazon EC2 in only one region, without limiting access to other services in any region. How can this company achieve these new security requirements while minimizing the administrative burden on the Operations team?

- A. Set up SAML-based authentication tied to an IAM role that has an AdministrativeAccess managed policy attached to it
- B. Attach a customer managed policy that denies access to Amazon EC2 in each region except for the one required.
- C. Create an IAM user for each Developer and add them to the developer IAM group that has the PowerUserAccess managed policy attached to it
- D. Attach a customer managed policy that allows the Developers access to Amazon EC2 only in the required region.
- E. Set up SAML-based authentication tied to an IAM role that has a PowerUserAccess managed policy and a customer managed policy that deny all the Developers access to any AWS services except AWS Service Catalog
- F. Within AWS Service Catalog, create a product containing only the EC2 resources in the approved region.
- G. Set up SAML-based authentication tied to an IAM role that has the PowerUserAccess managed policy attached to it
- H. Attach a customer managed policy that denies access to Amazon EC2 in each region except for the one required.

Answer: D

Explanation:

The tricks here are: - SAML for AD federation and authentication - PowerUserAccess vs AdministrativeAccess. (PowerUser has less privilege, which is the required one for developers). Admin, has more rights. The description of "PowerUser access" given by AWS is "Provides full access to AWS services and resources, but does not allow management of Users and groups."

NEW QUESTION 106

A company runs a Windows Server host in a public subnet that is configured to allow a team of administrators to connect over RDP to troubleshoot issues with hosts in a private subnet. The host must be available at all times outside of a scheduled maintenance window, and needs to receive the latest operating system updates within 3 days of release.

What should be done to manage the host with the LEAST amount of administrative effort?

- A. Run the host in a single-instance AWS Elastic Beanstalk environment
- B. Configure the environment with a custom AMI to use a hardened machine image from AWS Marketplace
- C. Apply system updates with AWS Systems Manager Patch Manager.
- D. Run the host on AWS WorkSpace

- E. Use Amazon WorkSpaces Application Manager (WAM) to harden the host
- F. Configure Windows automatic updates to occur every 3 days.
- G. Run the host in an Auto Scaling group with a minimum and maximum instance count of 1. Use a hardened machine image from AWS Marketplace
- H. Apply system updates with AWS Systems Manager Patch Manager.
- I. Run the host in AWS OpsWorks Stack
- J. Use a Chef recipe to harden the AMI during instance launch. Use an AWS Lambda scheduled event to run the Upgrade Operating System stack command to apply system updates.

Answer: B

NEW QUESTION 107

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

AWS Organizations has been used to create a new account and an AWS IAM user with administrator permissions for the Team Manager.

Which design meets these requirements?

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

Answer: D

NEW QUESTION 110

A Solutions Architect has been asked to look at a company's Amazon Redshift cluster, which has quickly become an integral part of its technology and supports key business processes. The Solutions Architect is to increase the reliability and availability of the cluster and provide options to ensure that if an issue arises, the cluster can either operate or be restored within four hours.

Which of the following solution options BEST addresses the business need in the most cost-effective manner?

- A. Ensure that the Amazon Redshift cluster has been set up to make use of Auto Scaling groups with the nodes in the cluster spread across multiple Availability Zones.
- B. Ensure that the Amazon Redshift cluster creation has been templated using AWS CloudFormation so it can easily be launched in another Availability Zone and data populated from the automated Redshift back-ups stored in Amazon S3.
- C. Use Amazon Kinesis Data Firehose to collect the data ahead of ingestion into Amazon Redshift and create clusters using AWS CloudFormation in another region and stream the data to both clusters.
- D. Create two identical Amazon Redshift clusters in different regions (one as the primary, one as the secondary). Use Amazon S3 cross-region replication from the primary to secondary. Use Amazon S3 cross-region replication from the primary to secondary region, which triggers an AWS Lambda function to populate the cluster in the secondary region.

Answer: B

Explanation:

https://aws.amazon.com/redshift/faqs/?nc1=h_ls Q: What happens to my data warehouse cluster availability and data durability if my data warehouse cluster's Availability Zone (AZ) has an outage? If your Amazon Redshift data warehouse cluster's Availability Zone becomes unavailable, you will not be able to use your cluster until power and network access to the AZ are restored. Your data warehouse cluster's data is preserved so you can start using your Amazon Redshift data warehouse as soon as the AZ becomes available again. In addition, you can also choose to restore any existing snapshots to a new AZ in the same Region. Amazon Redshift will restore your most frequently accessed data first so you can resume queries as quickly as possible.

FROM 37

NEW QUESTION 111

A company runs a video processing platform. Files are uploaded by users who connect to a web server, which stores them on an Amazon EFS share. This web server is running on a single Amazon EC2 instance. A different group of instances, running in an Auto Scaling group, scans the EFS share directory structure for new files to process and generates new videos (thumbnails, different resolution, compression, etc.) according to the instructions file, which is uploaded along with the video files. A different application running on a group of instances managed by an Auto Scaling group processes the video files and then deletes them from the EFS share. The results are stored in an S3 bucket. Links to the processed video files are emailed to the customer.

The company has recently discovered that as they add more instances to the Auto Scaling Group, many files are processed twice, so image processing speed is not improved. The maximum size of these video files is 2GB.

What should the Solutions Architect do to improve reliability and reduce the redundant processing of video files?

- A. Modify the web application to upload the video files directly to Amazon S3. Use Amazon CloudWatch Events to trigger an AWS Lambda function every time a file is uploaded, and have this Lambda function put a message into an Amazon SQS queue
- B. Modify the video processing application to read from SQS queue for new files and use the queue depth metric to scale instances in the video processing Auto Scaling group.
- C. Set up a cron job on the web server instance to synchronize the contents of the EFS share into Amazon S3. Trigger an AWS Lambda function every time a file is uploaded to process the video file and store the results in Amazon S3. Using Amazon CloudWatch Events trigger an Amazon SES job to send an email to the customer containing the link to the processed file.
- D. Rewrite the web application to run directly from Amazon S3 and use Amazon API Gateway to upload the video files to an S3 bucket
- E. Use an S3 trigger to run an AWS Lambda function each time a file is uploaded to process and store new video files in a different bucket
- F. Using CloudWatch Events, trigger an SES job to send an email to the customer containing the link to the processed file.

- G. Rewrite the web application to run from Amazon S3 and upload the video files to an S3 bucket
- H. Each time a new file is uploaded, trigger an AWS Lambda function to put a message in an SQS queue containing the link and the instruction
- I. Modify the video processing application to read from the SQS queue and the S3 bucket
- J. Use the queue depth metric to adjust the size of the Auto Scaling group for video processing instances.

Answer: A

NEW QUESTION 113

The Solutions Architect manages a serverless application that consists of multiple API gateways, AWS Lambda functions, Amazon S3 buckets, and Amazon DynamoDB tables. Customers say that a few application components slow while loading dynamic images, and some are timing out with the “504 Gateway Timeout” error. While troubleshooting the scenario, the Solutions Architect confirms that DynamoDB monitoring metrics are at acceptable levels. Which of the following steps would be optimal for debugging these application issues? (Choose two.)

- A. Parse HTTP logs in Amazon API Gateway for HTTP errors to determine the root cause of the errors.
- B. Parse Amazon CloudWatch Logs to determine processing times for requested images at specified intervals.
- C. Parse VPC Flow Logs to determine if there is packet loss between the Lambda function and S3.
- D. Parse AWS X-Ray traces and analyze HTTP methods to determine the root cause of the HTTP errors.
- E. Parse S3 access logs to determine if objects being accessed are from specific IP addresses to narrow the scope to geographic latency issues.

Answer: BD

Explanation:

Firstly “A 504 Gateway Timeout Error means your web server didn't receive a timely response from another server upstream when it attempted to load one of your web pages. Put simply, your web servers aren't communicating with each other fast enough”. This specific issue is addressed in the AWS article “Tracing, Logging and Monitoring an API Gateway API”. https://docs.amazonaws.cn/en_us/apigateway/latest/developerguide/monitoring_overview.html

NEW QUESTION 116

A company is moving a business-critical application onto AWS. It is a traditional three-tier web application using an Oracle database. Data must be encrypted in transit and at rest. The database hosts 12 TB of data. Network connectivity to the source Oracle database over the internet is allowed, and the company wants to reduce the operational costs by using AWS Managed Services where possible. All resources within the web and application tiers have been migrated. The database has a few tables and a simple schema using primary keys only; however, it contains many Binary Large Object (BLOB) fields. It was not possible to use the database's native replication tools because of licensing restrictions.

Which database migration solution will result in the LEAST amount of impact to the application's availability?

- A. Provision an Amazon RDS for Oracle instance
- B. Host the RDS database within a virtual private cloud (VPC) subnet with internet access, and set up the RDS database as an encrypted Read Replica of the source database
- C. Use SSL to encrypt the connection between the two databases
- D. Monitor the replication performance by watching the RDS ReplicaLag metric
- E. During the application maintenance window, shut down the on-premises database and switch over the application connection to the RDS instance when there is no more replication lag
- F. Promote the Read Replica into a standalone database instance.
- G. Provision an Amazon EC2 instance and install the same Oracle database software
- H. Create a backup of the source database using the supported tool
- I. During the application maintenance window, restore the backup into the Oracle database running in the EC2 instance
- J. Set up an Amazon RDS for Oracle instance, and create an import job between the database hosted in AWS
- K. Shut down the source database and switch over the database connections to the RDS instance when the job is complete.
- L. Use AWS DMS to load and replicate the dataset between the on-premises Oracle database and the replication instance hosted on AWS
- M. Provision an Amazon RDS for Oracle instance with Transparent Data Encryption (TDE) enabled and configure it as target for the replication instance
- N. Create a customer-managed AWS KMS master key to set it as the encryption key for the replication instance. Use AWS DMS tasks to load the data into the target RDS instance
- O. During the application maintenance window and after the load tasks reach the ongoing replication phase, switch the database connections to the new database.
- P. Create a compressed full database backup on the on-premises Oracle database during an application maintenance window
- Q. While the backup is being performed, provision a 10 Gbps AWS Direct Connect connection to increase the transfer speed of the database backup files to Amazon S3, and shorten the maintenance window period
- R. Use SSL/TLS to copy the files over the Direct Connect connection
- S. When the backup files are successfully copied, start the maintenance window, and use any of the Amazon RDS supported tools to import the data into a newly provisioned Amazon RDS for Oracle instance with encryption enabled
- T. Wait until the data is fully loaded and switch over the database connections to the new database
- . Delete the Direct Connect connection to cut unnecessary charges.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/apn/oracle-database-encryption-options-on-amazon-rds/>

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Appendix.Oracle.Options.AdvSecurity.htm> | (DMS in transit encryption)

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

NEW QUESTION 118

A Solutions Architect is migrating a 10 TB PostgreSQL database to Amazon RDS for PostgreSQL. The company's internet link is 50 MB with a VPN in the Amazon VPC, and the Solutions Architect needs to migrate the data and synchronize the changes before the cutover. The cutover must take place within an 8-day period.

What is the LEAST complex method of migrating the database securely and reliably?

- A. Order an AWS Snowball device and copy the database using the AWS Data Migration Agent
- B. When the database is available in Amazon S3, use AWS DMS to load it to Amazon RDS, and configure a job to synchronize changes before the cutover.
- C. Create an AWS DMS job to continuously replicate the data from on premises to AWS
- D. Cutover to Amazon RDS after the data is synchronized.
- E. Order an AWS Snowball device and copy a database dump to the device
- F. After the data has been copied to Amazon S3, import it to the Amazon RDS instance

- G. Set up log shipping over a VPN to synchronize changes before the cutover.
 H. Order an AWS Snowball device and copy the database by using the AWS Schema Conversion Tool. When the data is available in Amazon S3, use AWS DMS to load it to Amazon RDS, and configure a job to synchronize changes before the cutover.

Answer: B

NEW QUESTION 119

A company has several teams, and each team has their own Amazon RDS database that totals 100 TB. The company is building a data query platform for Business Intelligence Analysts to generate a weekly business report. The new system must run ad-hoc SQL queries. What is the MOST cost-effective solution?

- A. Create a new Amazon Redshift cluster. Create an AWS Glue ETL job to copy data from the RDS databases to the Amazon Redshift cluster. Use Amazon Redshift to run the query.
 B. Create an Amazon EMR cluster with enough core nodes. Run an Apache Spark job to copy data from the RDS databases to an Hadoop Distributed File System (HDFS). Use a local Apache Hive metastore to maintain the table definition. Use Spark SQL to run the query.
 C. Use an AWS Glue ETL job to copy all the RDS databases to a single Amazon Aurora PostgreSQL database. Run SQL queries on the Aurora PostgreSQL database.
 D. Use an AWS Glue crawler to crawl all the databases and create tables in the AWS Glue Data Catalog. Use an AWS Glue ETL Job to load data from the RDS databases to Amazon S3, and use Amazon Athena to run the queries.

Answer: C

NEW QUESTION 122

A company that provides wireless services needs a solution to store and analyze log files about user activities. Currently, log files are delivered daily to Amazon Linux on Amazon EC2 instance. A batch script is run once a day to aggregate data used for analysis by a third-party tool. The data pushed to the third-party tool is used to generate a visualization for end users. The batch script is cumbersome to maintain, and it takes several hours to deliver the ever-increasing data volumes to the third-party tool. The company wants to lower costs, and is open to considering a new tool that minimizes development effort and lowers administrative overhead. The company wants to build a more agile solution that can store and perform the analysis in near-real time, with minimal overhead. The solution needs to be cost effective and scalable to meet the company's end-user base growth. Which solution meets the company's requirements?

- A. Develop a Python script to failover the data from Amazon EC2 in real time and store the data in Amazon S3. Use a copy command to copy data from Amazon S3 to Amazon Redshift.
 B. Connect a business intelligence tool running on Amazon EC2 to Amazon Redshift and create the visualizations.
 C. Use an Amazon Kinesis agent running on an EC2 instance in an Auto Scaling group to collect and send the data to an Amazon Kinesis Data Firehose delivery stream.
 D. The Kinesis Data Firehose delivery stream will deliver the data directly to Amazon E.
 E. Use Kibana to visualize the data.
 F. Use an in-memory caching application running on an Amazon EBS-optimized EC2 instance to capture the log data in near real-time.
 G. Install an Amazon ES cluster on the same EC2 instance to store the log files as they are delivered to Amazon EC2 in near real-time.
 H. Install a Kibana plugin to create the visualizations.
 I. Use an Amazon Kinesis agent running on an EC2 instance to collect and send the data to an Amazon Kinesis Data Firehose delivery stream.
 J. The Kinesis Data Firehose delivery stream will deliver the data to Amazon S3. Use an AWS Lambda function to deliver the data from Amazon S3 to Amazon E.
 K. Use Kibana to visualize the data.

Answer: B

Explanation:

<https://docs.aws.amazon.com/firehose/latest/dev/writing-with-agents.html>

NEW QUESTION 126

AnyCompany has acquired numerous companies over the past few years. The CIO for AnyCompany would like to keep the resources for each acquired company separate. The CIO also would like to enforce a chargeback model where each company pays for the AWS services it uses. The Solutions Architect is tasked with designing an AWS architecture that allows AnyCompany to achieve the following:

- Implementing a detailed chargeback mechanism to ensure that each company pays for the resources it uses.
- AnyCompany can pay for AWS services for all its companies through a single invoice.
- Developers in each acquired company have access to resources in their company only.
- Developers in an acquired company should not be able to affect resources in their company only.
- A single identity store is used to authenticate Developers across all companies.

Which of the following approaches would meet these requirements? (Choose two.)

- A. Create a multi-account strategy with an account per company.
 B. Use consolidated billing to ensure that AnyCompany needs to pay a single bill only.
 C. Create a multi-account strategy with a virtual private cloud (VPC) for each company.
 D. Reduce impact across companies by not creating any VPC peering link.
 E. As everything is in a single account, there will be a single invoice.
 F. Use tagging to create a detailed bill for each company.
 G. Create IAM users for each Developer in the account to which they require access.
 H. Create policies that allow the users access to all resources in that account.
 I. Attach the policies to the IAM user.
 J. Create a federated identity store against the company's Active Directory.
 K. Create IAM roles with appropriate permissions and set the trust relationships with AWS and the identity store.
 L. Use AWS STS to grant users access based on the groups they belong to in the identity store.
 M. Create a multi-account strategy with an account per company.
 N. For billing purposes, use a tagging solution that uses a tag to identify the company that creates each resource.

Answer: AD

NEW QUESTION 130

A company is migrating to the cloud. It wants to evaluate the configurations of virtual machines in its existing data center environment to ensure that it can size new Amazon EC2 instances accurately. The company wants to collect metrics, such as CPU, memory, and disk utilization, and it needs an inventory of what processes are running on each instance. The company would also like to monitor network connections to map communications between servers. Which would enable the collection of this data MOST cost effectively?

- A. Use AWS Application Discovery Service and deploy the data collection agent to each virtual machine in the data center.
- B. Configure the Amazon CloudWatch agent on all servers within the local environment and publish metrics to Amazon CloudWatch Logs.
- C. Use AWS Application Discovery Service and enable agentless discovery in the existing virtualization environment.
- D. Enable AWS Application Discovery Service in the AWS Management Console and configure the corporate firewall to allow scans over a VPN.

Answer: A

NEW QUESTION 131

A Solutions Architect must create a cost-effective backup solution for a company's 500MB source code repository of proprietary and sensitive applications. The repository runs on Linux and backs up daily to tape. Tape backups are stored for 1 year. The current solutions are not meeting the company's needs because it is a manual process that is prone to error, expensive to maintain, and does not meet the need for a Recovery Point Objective (RPO) of 1 hour or Recovery Time Objective (RTO) of 2 hours. The new disaster recovery requirement is for backups to be stored offsite and to be able to restore a single file if needed. Which solution meets the customer's needs for RTO, RPO, and disaster recovery with the LEAST effort and expense?

- A. Replace local tapes with an AWS Storage Gateway virtual tape library to integrate with current backup software
- B. Run backups nightly and store the virtual tapes on Amazon S3 standard storage in US-EAST-1. Use cross-region replication to create a second copy in US-WEST-2. Use Amazon S3 lifecycle policies to perform automatic migration to Amazon Glacier and deletion of expired backups after 1 year?
- C. Configure the local source code repository to synchronize files to an AWS Storage Gateway file Amazon gateway to store backup copies in an Amazon S3 Standard bucket
- D. Enable versioning on the Amazon S3 bucket
- E. Create Amazon S3 lifecycle policies to automatically migrate old versions of objects to Amazon S3 Standard 0 Infrequent Access, then Amazon Glacier, then delete backups after 1 year.
- F. Replace the local source code repository storage with a Storage Gateway stored volume
- G. Change the default snapshot frequency to 1 hour
- H. Use Amazon S3 lifecycle policies to archive snapshots to Amazon Glacier and remove old snapshots after 1 year
- I. Use cross-region replication to create a copy of the snapshots in US-WEST-2.
- J. Replace the local source code repository storage with a Storage Gateway cached volume
- K. Create a snapshot schedule to take hourly snapshots
- L. Use an Amazon CloudWatch Events schedule expression rule to run on hourly AWS Lambda task to copy snapshots from US-EAST -1 to US-WEST-2.

Answer: B

Explanation:

<https://d1.awsstatic.com/whitepapers/aws-storage-gateway-file-gateway-for-hybrid-architectures.pdf>

NEW QUESTION 135

A company manages more than 200 separate internet-facing web applications. All of the applications are deployed to AWS in a single AWS Region. The fully qualified domain names (FQDNs) of all of the applications are made available through HTTPS using Application Load Balancers (ALBs). The ALBs are configured to use public SSL/TLS certificates.

A Solutions Architect needs to migrate the web applications to a multi-region architecture. All HTTPS services should continue to work without interruption. Which approach meets these requirements?

- A. Request a certificate for each FQDN using AWS KMS
- B. Associate the certificates with the ALBs in the primary AWS Region
- C. Enable cross-region availability in AWS KMS for the certificates and associate the certificates with the ALBs in the secondary AWS Region.
- D. Generate the key pairs and certificate requests for each FQDN using AWS KMS
- E. Associate the certificates with the ALBs in both the primary and secondary AWS Regions.
- F. Request a certificate for each FQDN using AWS Certificate Manager
- G. Associate the certificates with the ALBs in both the primary and secondary AWS Regions.
- H. Request certificates for each FQDN in both the primary and secondary AWS Regions using AWS Certificate Manager
- I. Associate the certificates with the corresponding ALBs in each AWS Region.

Answer: D

Explanation:

<https://docs.aws.amazon.com/acm/latest/userguide/acm-regions.html>

Certificates in ACM are regional resources. To use a certificate with Elastic Load Balancing for the same fully qualified domain name (FQDN) or set of FQDNs in more than one AWS region, you must request or import a certificate for each region. For certificates provided by ACM, this means you must revalidate each domain name in the certificate for each region. You cannot copy a certificate between regions.

NEW QUESTION 139

A company has more than 100 AWS accounts, with one VPC per account, that need outbound HTTPS connectivity to the internet. The current design contains one NAT gateway per Availability Zone (AZ) in each VPC. To reduce costs and obtain information about outbound traffic, management has asked for a new architecture for internet access.

Which solution will meet the current needs, and continue to grow as new accounts are provisioned, while reducing costs?

- A. Create a transit VPC across two AZs using a third-party routing appliance
- B. Create a VPN connection to each VPC
- C. Default route internet traffic to the transit VPC.
- D. Create multiple hosted-private AWS Direct Connect VIFs, one per account, each with a Direct Connect gateway
- E. Default route internet traffic back to an on-premises router to route to the internet.
- F. Create a central VPC for outbound internet traffic

- G. Use VPC peering to default route to a set of redundant NAT gateway in the central VPC.
- H. Create a proxy fleet in a central VPC account
- I. Create an AWS PrivateLink endpoint service in the central VPC
- J. Use PrivateLink interface for internet connectivity through the proxy fleet.

Answer: D

Explanation:

user proxy fleet over PrivateLink. As explained in this AWS website:
<https://aws.amazon.com/blogs/networking-and-content-delivery/how-to-use-aws-privatelink-to-secure-and-scale>

NEW QUESTION 141

An internal security audit of AWS resources within a company found that a number of Amazon EC2 instances running Microsoft Windows workloads were missing several important operating system-level patches. A Solutions Architect has been asked to fix existing patch deficiencies, and to develop a workflow to ensure that future patching requirements are identified and taken care of quickly. The Solutions Architect has decided to use AWS Systems Manager. It is important that EC2 instance reboots do not occur at the same time on all Windows workloads to meet organizational uptime requirements. Which workflow will meet these requirements in an automated manner?

- A. Add a Patch Group tag with a value of Windows Servers to all existing EC2 instances
- B. Ensure that all Windows EC2 instances are assigned this tag
- C. Associate the AWS-DefaultPatchBaseline to the Windows servers patch group
- D. Define an AWS Systems Manager maintenance window, conduct patching within it, and associate it with the Windows Servers patch group
- E. Register instances with the maintenance window using associated subnet ID
- F. Assign the AWS-RunPatchBaseline document as a task within each maintenance window.
- G. Add a Patch Group tag with a value of Windows Servers to all existing EC2 instances
- H. Ensure that all Windows EC2 instances are assigned this tag
- I. Associate the AWS-WindowsPatchBaseline document as a task associated with the Windows Servers patch group
- J. Create an Amazon CloudWatch Events rule configured to use a cron expression to schedule the execution of patching using the AWS Systems Manager run command
- K. Create an AWS Systems Manager State Manager document to define commands to be executed during patch execution.
- L. Add a Patch Group tag with a value of either Windows Servers1 or Windows Servers2 to all existing EC2 instances
- M. Ensure that all Windows EC2 instances are assigned this tag
- N. Associate the AWS-DefaultPatchBaseline with both Windows Servers patch groups
- O. Define two non-overlapping AWS Systems Manager maintenance windows, conduct patching within them, and associate each with a different patch group
- P. Register targets with specific maintenance windows using the Patch Group tag
- Q. Assign the AWS-RunPatchBaseline document as a task within each maintenance window.
- R. Add a Patch Group tag with a value of either Windows Servers1 or Windows Servers2 to all existing EC2 instances
- S. Ensure that all Windows EC2 instances are assigned this tag
- T. Associate the AWS-WindowsPatchBaseline with both Windows Servers patch groups
- U. Define two non-overlapping AWS Systems Manager maintenance windows, conduct patching within them, and associate each with a different patch group
- V. Assign the AWS-RunWindowsPatchBaseline document as a task within each maintenance window
- W. Create an AWS Systems Manager State Manager document to define commands to be executed during patch execution.

Answer: C

NEW QUESTION 146

A company operating a website on AWS requires high levels of scalability, availability and performance. The company is running a Ruby on Rails application on Amazon EC2. It has a data tier on MySQL 5.6 on Amazon EC2 using 16 TB of Amazon EBS storage. Amazon CloudFront is used to cache application content. The Operations team is reporting continuous and unexpected growth of EBS volumes assigned to the MySQL database. The Solutions Architect has been asked to design a highly scalable, highly available, and high-performing solution. Which solution is the MOST cost-effective at scale?

- A. Implement Multi-AZ and Auto Scaling for all EC2 instances in the current configuration
- B. Ensure that all EC2 instances are purchased as reserved instances
- C. Implement new elastic Amazon EBS volumes for the data tier.
- D. Design and implement the Docker-based containerized solution for the application using Amazon EC2
- E. Migrate to an Amazon Aurora MySQL Multi-AZ cluster
- F. Implement storage checks for Aurora MySQL storage utilization and an AWS Lambda function to grow the Aurora MySQL storage, as necessary
- G. Ensure that Multi-AZ architectures are implemented.
- H. Ensure that EC2 instances are right-sized and behind an Elastic Load Balancing load balancer. Implement Auto Scaling with EC2 instances
- I. Ensure that the reserved instances are purchased for fixed capacity and that Auto Scaling instances run on demand
- J. Migrate to an Amazon Aurora MySQL Multi-AZ cluster
- K. Ensure that Multi-AZ architectures are implemented.
- L. Ensure that EC2 instances are right-sized and behind an Elastic Load Balancing
- M. Implement Auto Scaling with EC2 instances
- N. Ensure that Reserved instances are purchased for fixed capacity and that Auto Scaling instances run on demand
- O. Migrate to an Amazon Aurora MySQL Multi-AZ cluster
- P. Implement storage checks for Aurora MySQL storage utilization and an AWS Lambda function to grow Aurora MySQL storage, as necessary
- Q. Ensure Multi-AZ architectures are implemented.

Answer: C

NEW QUESTION 150

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

- A. Receive the orders in an Amazon EC2-hosted database and use EC2 instances to process them.

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

Answer: B

NEW QUESTION 151

A company has a large on-premises Apache Hadoop cluster with a 20 PB HDFS database. The cluster is growing every quarter by roughly 200 instances and 1 PB. The company's goals are to enable resiliency for its Hadoop data, limit the impact of losing cluster nodes, and significantly reduce costs. The current cluster runs 24/7 and supports a variety of analysis workloads, including interactive queries and batch processing.

Which solution would meet these requirements with the LEAST expense and down time?

- A. Use AWS Snowmobile to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster initially sized to handle the interactive workload based on historical data from the on-premises cluster
- B. Store the data on EMRFS
- C. Minimize costs using Reserved Instances for master and core nodes and Spot Instances for task nodes, and auto scale task nodes based on Amazon CloudWatch metric
- D. Create job-specific, optimized clusters for batch workloads that are similarly optimized.
- E. Use AWS Snowmobile to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster of similar size and configuration to the current cluster
- F. Store the data on EMRFS
- G. Minimize costs by using Reserved Instance
- H. As the workload grows each quarter, purchase additional Reserved Instances and add to the cluster.
- I. Use AWS Snowball to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster initially sized to handle the interactive workloads based on historical data from the on-premises cluster
- J. Store the data on EMRFS
- K. Minimize costs using Reserved Instances for master and core nodes and Spot Instances for task nodes, and auto scale task nodes based on Amazon CloudWatch metric
- L. Create job-specific, optimized clusters for batch workloads that are similarly optimized.
- M. Use AWS Direct Connect to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster initially sized to handle the interactive workload based on historical data from the on-premises cluster
- N. Store the data on EMRFS
- O. Minimize costs using Reserved Instances for master and core nodes and Spot Instances for task nodes, and auto scale task nodes based on Amazon CloudWatch metric
- P. Create job-specific, optimized clusters for batch workloads that are similarly optimized.

Answer: A

Explanation:

Q: How should I choose between Snowmobile and Snowball?

To migrate large datasets of 10PB or more in a single location, you should use Snowmobile. For datasets less than 10PB or distributed in multiple locations, you should use Snowball. In addition, you should evaluate the amount of available bandwidth in your network backbone. If you have a high speed backbone with hundreds of Gb/s of spare throughput, then you can use Snowmobile to migrate the large datasets all at once. If you have limited bandwidth on your backbone, you should consider using multiple Snowballs to migrate the data incrementally.

NEW QUESTION 155

A Solutions Architect must migrate an existing on-premises web application with 70 TB of static files supporting a public open-data initiative. The architect wants to upgrade to the latest version of the host operating system as part of the migration effort.

Which is the FASTEST and MOST cost-effective way to perform the migration?

- A. Run a physical-to-virtual conversion on the application server
- B. Transfer the server image over the internet, and transfer the static data to Amazon S3.
- C. Run a physical-to-virtual conversion on the application server
- D. Transfer the server image over AWS Direct Connect, and transfer the static data to Amazon S3.
- E. Re-platform the server to Amazon EC2, and use AWS Snowball to transfer the static data to Amazon S3.
- F. Re-platform the server by using the AWS Server Migration Service to move the code and data to a new Amazon EC2 instance.

Answer: C

NEW QUESTION 157

A company has a website that enables users to upload videos. Company policy states the uploaded videos must be analyzed for restricted content. An uploaded video is placed in Amazon S3, and a message is pushed to an Amazon SQS queue with the video's location. A backend application pulls this location from Amazon SQS and analyzes the video.

The video analysis is compute-intensive and occurs sporadically during the day. The website scales with demand. The video analysis application runs on a fixed number of instances. Peak demand occurs during the holidays, so the company must add instances to the application during this time. All instances used are currently on-demand Amazon EC2 T2 instances. The company wants to reduce the cost of the current solution.

Which of the following solutions is MOST cost-effective?

- A. Keep the website on T2 instances. Determine the minimum number of website instances required during off-peak times and use Spot Instances to cover them while using Reserved Instances to cover peak demand. Use Amazon EC2 R4 and Amazon EC2 R5 Reserved Instances in an Auto Scaling group for the video analysis application.
- B. Keep the website on T2 instances. Determine the minimum number of website instances required during off-peak times and use Reserved Instances to cover them while using On-Demand Instances to cover peak demand. Use Spot Fleet for the video analysis application comprised of Amazon EC2 C4 and Amazon EC2 C5 Spot Instances.
- C. Migrate the website to AWS Elastic Beanstalk and Amazon EC2 C4 instances. Determine the minimum number of website instances required during off-peak times and use On-Demand instances to cover them while using Spot capacity to cover peak demand. Use Spot Fleet for the video analysis application comprised of C4 and Amazon EC2 C5 instances.
- D. Migrate the website to AWS Elastic Beanstalk and Amazon EC2 R4 instances. Determine the minimum number of website instances required during off-peak times and use Reserved Instances to cover them while using On-Demand Instances to cover peak demand. Use Spot Fleet for the video analysis application comprised of R4 and Amazon EC2 R5 instances.

Answer: B

NEW QUESTION 159

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

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

Answer: B

NEW QUESTION 160

The CISO of a large enterprise with multiple IT departments, each with its own AWS account, wants one central place where AWS permissions for users can be managed and users authentication credentials can be synchronized with the company's existing on-premises solution. Which solution will meet the CISO's requirements?

- A. Define AWS IAM roles based on the functional responsibilities of the users in a central account
- B. Create a SAML-based identity management provider
- C. Map users in the on-premises groups to IAM role
- D. Establish trust relationships between the other accounts and the central account.
- E. Deploy a common set of AWS IAM users, groups, roles, and policies in all of the AWS accounts using AWS Organization
- F. Implement federation between the on-premises identity provider and the AWS accounts.
- G. Use AWS Organizations in a centralized account to define service control policies (SCPs). Create a SAML-based identity management provider in each account and map users in the on-premises groups to AWS IAM roles.
- H. Perform a thorough analysis of the user base and create AWS IAM users accounts that have the necessary permissions
- I. Set up a process to provision and de-provision accounts based on data in the on-premises solution.

Answer: A

Explanation:

https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

NEW QUESTION 163

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

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

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

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

Answer: D

NEW QUESTION 164

A financial services company is moving to AWS and wants to enable Developers to experiment and innovate while preventing access to production applications. The company has the following requirements:

- Production workloads cannot be directly connected to the internet
- All workloads must be restricted to the us-west-2 and eu-central-1 Regions
- Notification should be sent when Developer sandboxes exceed \$500 in AWS spending monthly

Which combination of actions needs to be taken to create a multi-account structure that meets the company's requirements? (Select THREE)

- A. Create accounts for each production workload within an organization in AWS Organizations Place the production accounts within an organizational unit (OU) For each account delete the default VPC Create an SCP with a Deny rule for the attach an internet gateway and create a default VPC actions Attach the SCP to the OU for the production accounts

- B. Create accounts for each production workload within an organization in AWS Organizations Place the production accounts within an organizational unit (OU) Create an SCP with a Deny rule on the attach an internet gateway action Create an SCP with a Deny rule to prevent use of the default VPC Attach the SCPs to the OU for the production accounts
- C. Create a SCP containing a Deny Effect for cloudfront". lam:*, route53* and support* with a StringNotEquals condition on an aws RequestedRegion condition key with us-west-2 and eu-central-1 values Attach the SCP to the organization's root.
- D. Create an IAM permission boundary containing a Deny Effect for cloudfront'. lam * route53' and support" with a StringNotEquals condition on an aws RequestedRegion condition key with us-west 2 and eu-central-1 values Attach the permission boundary to an IAM group containing the development and production users.
- E. Create accounts for each development workload within an organization m AWS Organizations Place the development accounts within an organizational unit (OU) Create a custom AWS Config rule to deactivate all (AM users when an account's monthly bill exceeds \$500.
- F. Create accounts for each development workload within an organization in AWS Organizations Place the development accounts within an organizational unit (OU) Create a budget within AWS Budgets for each development account to monitor and report on monthly spending exceeding \$500.

Answer: ABD

NEW QUESTION 169

A company is running a web application with On-Demand Amazon EC2 instances in Auto Scaling groups that scale dynamically based on custom metrics After extensive testing the company determines that the m5 2xlarge instance size is optimal for the workload Application data is stored in db r4 4xlarge Amazon RDS instances that are confirmed to be optimal The traffic to the web application spikes randomly during the day What other cost-optimization methods should the company implement to further reduce costs without impacting the reliability of the application?

- A. Double the instance count in the Auto Scaling groups and reduce the instance size to m5 large
- B. Reserve capacity for the RDS database and the minimum number of EC2 instances that are constantly running
- C. Reduce the RDS instance size to db r4 xlarge and add five equivalents sized read replicas to provide reliability
- D. Reserve capacity for all EC2 instances and leverage Spot Instance pricing for the RDS database

Answer: B

NEW QUESTION 172

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

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

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

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

Answer: D

NEW QUESTION 176

A company has an application behind a load balancer with enough Amazon EC2 instances to satisfy peak demand. Scripts and third-party deployment solutions are used to configure EC2 instances when demand increases or an instance fails. The team must periodically evaluate the utilization of the instance types to ensure that the correct sizes are deployed.

How can this workload be optimized to meet these requirements?

- A. Use CloudFormer` to create AWS CloudFormation stacks from the current resource
- B. Deploy that stack by using AWS CloudFormation in the same regio
- C. Use Amazon CloudWatch alarms to send notifications about underutilized resources to provide cost-savings suggestions.
- D. Create an Auto Scaling group to scale the instances, and use AWS CodeDeploy to perform the configuratio
- E. Change from a load balancer to an Application Load Balance
- F. Purchase a third-party product that provides suggestions for cost savings on AWS resources.
- G. Deploy the application by using AWS Elastic Beanstalk with default option
- H. Register for an AWS Support Developer pla
- I. Review the instance usage for the application by using Amazon CloudWatch, and identify less expensive instances that can handle the loa
- J. Hold monthly meetings to review new instance types and determine whether Reserved instances should be purchased.
- K. Deploy the application as a Docker image by using Amazon EC
- L. Set up Amazon EC2 Auto Scaling and Amazon ECS scalin
- M. Register for AWS Business Support and use Trusted Advisor checks to provide suggestions on cost savings.

Answer: D

NEW QUESTION 179

A company wants to replace its call system with a solution built using AWS managed services. The company call center would like the solution to receive calls,

create contact flows, and scale to handle growth projections. The call center would also like the solution to use deep learning capabilities to recognize the intent of the callers and handle basic tasks, reducing the need to speak an agent. The solution should also be able to query business applications and provide relevant information back to calls as requested.

Which services should the Solution Architect use to build this solution? (Choose three.)

- A. Amazon Rekognition to identify who is calling.
- B. Amazon Connect to create a cloud-based contact center.
- C. Amazon Alexa for Business to build conversational interface.
- D. AWS Lambda to integrate with internal systems.
- E. Amazon Lex to recognize the intent of the caller.
- F. Amazon SQS to add incoming callers to a queue.

Answer: BDE

NEW QUESTION 181

A company currently uses a single 1 Gbps AWS Direct Connect connection to establish connectivity between an AWS Region and its data center. The company has five Amazon VPCs, all of which are connected to the data center using the same Direct Connect connection. The Network team is worried about the single point of failure and is interested in improving the redundancy of the connections to AWS while keeping costs to a minimum.

Which solution would improve the redundancy of the connection to AWS while meeting the cost requirements?

- A. Provision another 1 Gbps Direct Connect connection and create new VIFs to each of the VPCs. Configure the VIFs in a load balancing fashion using BGP.
- B. Set up VPN tunnels from the data center to each VPC.
- C. Terminate each VPN tunnel at the virtual private gateway (VGW) of the respective VPC and set up BGP for route management.
- D. Set up a new point-to-point Multiprotocol Label Switching (MPLS) connection to the AWS Region that's being used.
- E. Configure BGP to use this new circuit as passive, so that no traffic flows through this unless the AWS Direct Connect fails.
- F. Create a public VIF on the Direct Connect connection and set up a VPN tunnel which will terminate on the virtual private gateway (VGW) of the respective VPC using the public VIF.
- G. Use BGP to handle the failover to the VPN connection.

Answer: B

NEW QUESTION 182

A large company is migrating its entire IT portfolio to AWS. Each business unit in the company has a standalone AWS account that supports both development and test environments. New accounts to support production workloads will be needed soon.

The Finance department requires a centralized method for payment but must maintain visibility into each group's spending to allocate costs.

The Security team requires a centralized mechanism to control IAM usage in all the company's accounts. What combination of the following options meet the company's needs with LEAST effort? (Choose two.)

- A. Use a collection of parameterized AWS CloudFormation templates defining common IAM permissions that are launched into each account.
- B. Require all new and existing accounts to launch the appropriate stacks to enforce the least privilege model.
- C. Use AWS Organizations to create a new organization from a chosen payer account and define an organizational unit hierarchy.
- D. Invite the existing accounts to join the organization and create new accounts using Organizations.
- E. Require each business unit to use its own AWS account.
- F. Tag each AWS account appropriately and enable Cost Explorer to administer chargebacks.
- G. Enable all features of AWS Organizations and establish appropriate service control policies that filter IAM permissions for sub-accounts.
- H. Consolidate all of the company's AWS accounts into a single AWS account.
- I. Use tags for billing purposes and IAM's Access Advice feature to enforce the least privilege model.

Answer: BD

NEW QUESTION 186

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual SAP-C01 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the SAP-C01 Product From:

<https://www.2passeasy.com/dumps/SAP-C01/>

Money Back Guarantee

SAP-C01 Practice Exam Features:

- * SAP-C01 Questions and Answers Updated Frequently
- * SAP-C01 Practice Questions Verified by Expert Senior Certified Staff
- * SAP-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * SAP-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year