Exam Questions SAP-C01

AWS Certified Solutions Architect- Professional

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

NEW QUESTION 1
A Solutions Architect is building a containerized NET Core application that will run in AWS Fargate The backend of the application requires Microsoft SQL Server with high availability All tiers of the application must be highly available The credentials used for the connection string to SQL Server should not be stored on disk within the .NET Core front-end containers.
Which strategies should the Solutions Architect use to meet these requirements'?

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

Answer: C

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 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
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 scalable way?

A. Store the data in a single Amazon S3 bucke
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 bucke
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 typ
E. Users get credentials to access the data by using AssumeRole from their business unit's AWS accoun
F. Users can then use those credentials with an S3 client.
G. Store the data in a single Amazon S3 bucke
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 typ
I. Determine the business unit with the AWS account accessing the bucket and the job type with a prefix in the IAM user's nam
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 bucket
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 applicatio
M. The application uses the business unit and job type information in the IdP to control what users can upload and download through the applicatio
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 isintegrated 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 acces
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 4
A company wants to follow its website on AWS using serverless architecture design patterns for global customers. The company has outlined its requirements as follow:

> The website should be responsive.

> The website should offer minimal latency.

> The website should be highly available.

> Users should be able to authenticate through social identity providers such as Google, Facebook, and Amazon.

> There should be baseline DDoS protections for spikes in traffic.

How can the design requirements be met?

A. Use Amazon CloudFront with Amazon ECS for hosting the websit
B. Use AWS Secrets Manager for provide user management and authentication function
C. Use ECS Docker containers to build an API.
D. Use Amazon Route 53 latency routing with an Application Load Balancer and AWS Fargate in different regions for hosting the websit
E. use Amazon Cognito to provide user management and authentication function
F. Use Amazon EKS containers.
G. Use Amazon CloudFront with Amazon S3 for hosting static web resource
H. Use Amazon Cognito to provide user management authentication function
I. Use Amazon API Gateway with AWS Lambda to build an API.
J. Use AWS Direct Connect with Amazon CloudFront and Amazon S3 for hosting static web resource.Use Amazon Cognito to provide user management authentication function
K. Use AWS Lambda to build an API.

Answer: C

NEW QUESTION 5
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 traffi
B. Open the network ACL and prevent administration from making future changes through IAM.
C. The fault is in the third-party environmen
D. Contact the third party that provides the maps and request a fix that will provide better uptime.
E. One NAT instance has become overloade
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 faile
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 6
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 7
A company currently uses Amazon EBS and Amazon RDS for storage purposes. The company intends to use a pilot light approach for disaster recovery in a different AWS Region. The company has an RTO of 6 hours and an RPO of 24 hours.
Which solution would achieve the requirements with MINIMAL cost?

A. Use AWS Lambda to create daily EBS and RDS snapshots, and copy them to the disaster recovery regio
B. Use Amazon Route 53 with active-passive failover configuratio
C. Use Amazon EC2 in an Auto Scaling group with the capacity set to 0 in the disaster recovery region.
D. Use AWS Lambda to create daily EBS and RDS snapshots, and copy them to the disaster recovery regio
E. Use Amazon Route 53 with active-active failover configuratio
F. Use Amazon EC2 in an AutoScaling group configured in the same way as in the primary region.
G. Use Amazon ECS to handle long-running tasks to create daily EBS and RDS snapshots, and copy to the disaster recovery regio
H. Use Amazon Route 53 with active-passive failover configuratio
I. Use Amazon EC2 in an Auto Scaling group with the capacity set to 0 in the disaster recovery region.
J. Use EBS and RDS cross-region snapshot copy capability to create snapshots in the disaster recovery regio

K. Use Amazon Route 53 with active-active failover configuratio
L. Use Amazon EC2 in an Auto Scaling group with the capacity set to 0 in the disaster recovery region.

Answer: A

Explanation:
https://docs.aws.amazon.com/AmazonECS/latest/developerguide/scheduling_tasks.html

NEW QUESTION 8
A photo-sharing and publishing company receives 10,000 to 150,000 images daily. The company receives the images from multiple suppliers and users registered with the service. The company is moving to AWS and wants to enrich the existing metadata by adding data using Amazon Rekognition.
The following is an example of the additional data:

```
list celebrities [name of the personality] wearing [color] looking [happy, sad] near [location example Eiffel Tower in Paris]
```

As part of the cloud migration program, the company uploaded existing image data to Amazon S3 and told users to upload images directly to Amazon S3.
What should the Solutions Architect do to support these requirements?

A. Trigger AWS Lambda based on an S3 event notification to create additional metadata using Amazon Rekognitio
B. Use Amazon DynamoDB to store the metadata and Amazon ES to create an inde
C. Use a web front-end to provide search capabilities backed by Amazon ES.
D. Use Amazon Kinesis to stream data based on an S3 even
E. Use an application running in Amazon EC2 to extract metadata from the image
F. Then store the data on Amazon DynamoDB and Amazon CloudSearch and create an inde
G. Use a web front-end with search capabilities backed by CloudSearch.
H. Start an Amazon SQS queue based on S3 event notification
I. Then have Amazon SQS send the metadata information to Amazon DynamoD
J. An application running on Amazon EC2 extracts data from Amazon Rekognition using the API and adds data to DynamoDB and Amazon E
K. Use a web front-end to provide search capabilities backed by Amazon ES.
L. Trigger AWS Lambda based on an S3 event notification to create additional metadata using Amazon Rekognitio
M. Use Amazon RDS MySQL Multi-AZ to store the metadata information and use Lambda to create an inde
N. Use a web front-end with search capabilities backed by Lambda.

Answer: A

Explanation:
https://github.com/aws-samples/lambda-refarch-imagerecognition

NEW QUESTION 9
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 functio
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 cod
D. Rollback if Amazon CloudWatch alarms are triggered.
E. Refactor the AWS CLI scripts into a single script that deploys the new Lambda versio
F. When deployment is completed, the script tests execut
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 versio
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-serverle

NEW QUESTION 10
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 instanc
B. Configure an IAM role to allow the start and stop of the Elastic Beanstalk environment and attach it to the instanc
C. Create scripts on the instance to start and stop the Elastic Beanstalk environmen
D. Configure cron jobs on the instance to execute the scripts.
E. Develop AWS Lambda functions to start and stop the Elastic Beanstalk environmen
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 environmen
I. Create a role for Step Functionsto allow it to start and stop the Elastic Beanstalk environmen
J. Invoke Step Functions daily.
K. Configure a time-based Auto Scaling grou
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 dat
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 10
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 Consol
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 applicatio
D. Write shell scripts that implement the rest of the steps in the runboo
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 applicatio
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 13
A company would like to implement a serverless application by using Amazon API Gateway, AWS Lambda and Amazon DynamoDB. They deployed a proof of concept and stated that the average response time is greater than what their upstream services can accept Amazon CloudWatch metrics did not indicate any issues with DynamoDB but showed that some Lambda functions were hitting their timeout.
Which of the following actions should the Solutions Architect consider to improve performance? (Choose two.)

A. Configure the AWS Lambda function to reuse containers to avoid unnecessary startup time.
B. Increase the amount of memory and adjust the timeout on the Lambda functio
C. Complete performance testing to identify the ideal memory and timeout configuration for the Lambda function.
D. Create an Amazon ElastiCache cluster running Memcached, and configure the Lambda function for VPC integration with access to the Amazon ElastiCache cluster.
E. Enable API cache on the appropriate stage in Amazon API Gateway, and override the TTL for individual methods that require a lower TTL than the entire stage.
F. Increase the amount of CPU, and adjust the timeout on the Lambda functio
G. Complete performance testing to identify the ideal CPU and timeout configuration for the Lambda function.

Answer: BD

Explanation:
https://lumigo.io/blog/aws-lambda-timeout-best-practices/

NEW QUESTION 17
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 20
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 21
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 backing.
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 functio
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 queu
D. Extract the core processing code from the existing application and update it to pull items from Amazon SQS queu
E. Extract the core processing code from the existing application and update it to pull items from Amazon SQS instead of an in-memory queu
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 RD
H. Configure Auto Scaling for the DynamoDB tabl
I. Deploy the application within an Auto Scaling group with a scaling policy based on CPU utilizatio
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 queu
L. Build a Docker container image for the applicatio
M. Create an Amazon ECS task definition that includes the application container and a separate container to host Redi
N. Deploy the new task definition as an ECS service using AWS Fargate and enable Auto Scaling.

Answer: B

NEW QUESTION 23
During an audit a Security team discovered that a Development team was putting IAM user secret access keys in their code and then committing it to an AWS CodeCommit repository The Security team wants to automatically find and remediate instances of this security vulnerability
Which solution will ensure that the credentials are appropriately secured automatically?

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

Answer: C

NEW QUESTION 26
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 27
A company is operating a large customer service call center, and stores and processes call recordings with a custom application Approximately 2% of the call recording are transcribed by an offshore team for quality assurance purposes. These recordings take days. The company uses Linux servers for processing the call recording and managing the transcription queue. There is also a web application for the quality assurance staff to review and score call recordings.
The company plans to migrate the system to AWS to reduce storage costs and the time required to transcribe calls.
Which set of actions should be taken to meet the company's objectives?

A. Upload the call recording to Amazon S3 from the call cente
B. Set up an S3 lifecycle policy to move the call recordings to Amazon S3 Glacier after 90 day
C. Use an AWS Lambda trigger to transcribe the call recordings with Amazon Transcrib
D. Use Amazon S3, Amazon API Gateway and Lambda to host the review and scoring application.
E. Upload the call recordings to Amazon S3 from the call cente
F. Set up an S3 lifecycle policy to move the call recordings to Amazon S3 Glacier after 90 day
G. Use an AWS Lambda trigger to transcribe the call recordings with Amazon Mechanical trun
H. Use Amazon EC2 instances in an Auto Scaling group behind an Application Balancer to host the review and scoring application.
I. Use Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer to host the review and scoring applicatio
J. Upload the call recordings to this application from the call center and store them on an Amazon EFS mount poin
K. Use AWS Backup to archive the call recording after 90 day
L. Transcribe the call recordings with Amazon Transcribe.
M. Upload the call recording to Amazon S3 from the call center and put the object key in an Amazon SQS queue
N. Set up an S3 lifecycle policy to move the call recordings to Amazon S3 Glacier after 90 day
O. Use Amazon EC2 instances in the queue as the scaling metri
P. Use Amazon S3, Amazon API Gateway, and AWS Lambda to host the review and scoring application.

Answer: B

NEW QUESTION 32
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 33
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 36
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 41
A company is running an email application across multiple AWS Regions. The company uses Ohio (us-east-2) as the primary Region and Northern Virginia (us-east-1) as the Disaster Recovery (DR) Region. The data is continuously replicated from the primary Region to the DR Region by a single instance on the public subnet in both Regions. The replication messages between the Regions have a significant backlog during certain times of the day. The backlog clears on its own after a short time, but it affects the application's RPO.
Which of the following solutions should help remediate this performance problem? (Select TWO)

A. Increase the size of the instances.
B. Have the instance in the primary Region write the data to an Amazon SQS queue in the primary Region instead, and have the instance in the DR Region poll from this queue.
C. Use multiple instances on the primary and DR Regions to send and receive the replication data.
D. Change the DR Region to Oregon (us-west-2) instead of the current DR Region.
E. Attach an additional elastic network interface to each of the instances in both Regions and set up load balancing between the network interfaces.

Answer: AC

NEW QUESTION 44
A Solutions Architect wants to make sure that only AWS users or roles with suitable permissions can access a new Amazon API Gateway endpoint The Solutions Architect wants an end-to-end view of each request to analyze the latency of the request and create service maps
How can the Solutions Architect design the API Gateway access control and perform request inspections?

A. For the API Gateway method set the authorization to AWSJAM Then, give the I AM user or role execute-api Invoke permission on the REST API resource Enable the API caller to sign requests with AWS Signature when accessing the endpoint Use AWS X-Roy to trace and analyze user requests to API Gateway
B. For the API Gateway resource set CORS to enabled and only return the company's domain mAccess-Control-Allow-Origin headers Then give the IAM user or role execute-api Invoke permission on the REST API resource Use Amazon CloudWatch to trace and analyze user requests to API Gateway
C. Create an AWS Lambda function as the custom authorizer ask the API client to pass the key and secret when making the call and then use Lambda to validate the key'secret pair against the IAM system Use AWS X-Ray to trace and analyze user requests to API Gateway
D. Create a client certificate for API Gateway Distribute the certificate to the AWS users and roles that need to access the endpoint Enable the API caller to pass the client certificate when accessing the endpoint Use Amazon CloudWatch to trace and analyze user requests to API Gateway.

Answer: D

NEW QUESTION 46
A company runs a public-facing application that uses a Java-based web sen/ice via a RESTful API It is hosted on Apache Tomcat on a single server in a data center that runs consistently at 30% CPU utilization Use of the API is expected to increase by 10 times with a new product launch The business wants to migrate the application to AWS with no disruption and needs it to scale to meet demand
The company has already decided to use Amazon Route 53 and CNAME records lo redirect traffic How can these requirements be met with the LEAST amount of effort?

A. Use AWS Elastic Beanstalk to deploy the Java web service and enable Auto Scaling Then switch the application to use the new web service
B. Lift and shift the Apache server to the cloud using AWS SMS Then switch the application to direct web service traffic to the new instance
C. Create a Docker image and migrate the image to Amazon ECS Then change the application code to direct web service queries to the ECS container
D. Modify the application to call the web service via Amazon API Gateway Then create a new AWS Lambda Java function to run the Java web service code After testing change API Gateway to use the Lambda function

Answer: A

NEW QUESTION 50
A company runs a three-tier application in AWS. Users report that the application performance can vary greatly depending on the time of day and functionality being accessed.
The application includes the following components:

> Eight t2.large front-end web servers that serve static content and proxy dynamic content from the application tier.
>

> Four t2.large application servers.

> One db.m4.large Amazon RDS MySQL Multi-AZ DB instance.

Operations has determined that the web and application tiers are network constrained.
Which of the following should cost effective improve application performance? (Choose two.)

A. Replace web and app tiers with t2.xlarge instances
B. Use AWS Auto Scaling and m4.large instances for the web and application tiers
C. Convert the MySQL RDS instance to a self-managed MySQL cluster on Amazon EC2
D. Create an Amazon CloudFront distribution to cache content
E. Increase the size of the Amazon RDS instance to db.m4.xlarge

Answer: BD

Explanation:
https://aws.amazon.com/ec2/instance-types/

NEW QUESTION 54
A company is implementing a multi-account strategy; however, the Management team has expressed concerns that services like DNS may become overly complex. The company needs a solution that allows private DNS to be shared among virtual private clouds (VPCs) in different accounts. The company will have approximately 50 accounts in total.
What solution would create the LEAST complex DNS architecture and ensure that each VPC can resolve all AWS resources?

A. Create a shared services VPC in a central account, and create a VPC peering connection from the shared services VPC to each of the VPCs in the other account
B. Within Amazon Route 53, create a privately hosted zone in the shared services VPC and resource record sets for the domain and subdomains.Programmatically associate other VPCs with the hosted zone.
C. Create a VPC peering connection among the VPCs in all account
D. Set the VPC attributes enableDnsHostnames and enableDnsSupport to "true" for each VP
E. Create an Amazon Route 53 private zone for each VP
F. Create resource record sets for the domain and subdomain
G. Programmatically associate the hosted zones in each VPC with the other VPCs.
H. Create a shared services VPC in a central accoun
I. Create a VPC peering connection from the VPCs in other accounts to the shared services VP
J. Create an Amazon Route 53 privately hosted zone in the shared services VPC with resource record sets for the domain and subdomain
K. Allow UDP and TCP port 53 over the VPC peering connections.
L. Set the VPC attributes enableDnsHostnames and enableDnsSupport to "false" in every VP
M. Create an AWS Direct Connect connection with a private virtual interfac
N. Allow UDP and TCP port 53 over the virtual interfac
O. Use the on-premises DNS servers to resolve the IP addresses in each VPC on AWS.

Answer: A

Explanation:
https://aws.amazon.com/blogs/networking-and-content-delivery/centralized-dns-management-of-hybrid-cloud-w

NEW QUESTION 55
A company has multiple AWS accounts hosting IT applications. An Amazon CloudWatch Logs agent is installed on all Amazon EC2 instances. The company wants to aggregate all security events in a centralized AWS account dedicated to log storage.
Security Administrators need to perform near-real-time gathering and correlating of events across multiple AWS accounts.
Which solution satisfies these requirements?

A. Create a Log Audit IAM role in each application AWS account with permissions to view CloudWatch Logs, configure an AWS Lambda function to assume the Log Audit role, and perform an hourly export of CloudWatch Logs data to an Amazon S3 bucket in the logging AWS account.
B. Configure CloudWatch Logs streams in each application AWS account to forward events to CloudWatch Logs in the logging AWS accoun
C. In the logging AWS account, subscribe an Amazon Kinesis Data Firehose stream to Amazon CloudWatch Events, and use the stream to persist log data in Amazon S3.
D. Create Amazon Kinesis Data Streams in the logging account, subscribe the stream to CloudWatch Logs streams in each application AWS account, configure an Amazon Kinesis Data Firehose delivery stream with the Data Streams as its source, and persist the log data in an Amazon S3 bucket inside the logging AWS account.
E. Configure CloudWatch Logs agents to publish data to an Amazon Kinesis Data Firehose stream in the logging AWS account, use an AWS Lambda function to read messages from the stream and push messages to Data Firehose, and persist the data in Amazon S3.

Answer: C

Explanation:
The solution uses Amazon Kinesis Data Streams and a log destination to set up an endpoint in the logging account to receive streamed logs and uses Amazon Kinesis Data Firehose to deliver log data to the Amazon Simple Storage Solution (S3) bucket. Application accounts will subscribe to stream all (or part) of their Amazon CloudWatch logs to a defined destination in the logging account via subscription filters. https://aws.amazon.com/blogs/architecture/central-logging-in-multi-account-environments/

NEW QUESTION 60
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 ECs 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 tea
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 teamonly.
F. Set up SAML federation for all accounts in AW
G. Configure SAML so that it checks for the service API call before authenticating the use
H. Block SAML from authenticating API calls if anyone other than the Security team accesses these instances.

Answer: B

NEW QUESTION 65
A company runs a dynamic mission-critical web application that has an SLA of 99.99%. Global application users access the application 24/7. The application is currently hosted on premises and routinely fails to meet its SLA, especially when millions of users access the application concurrently. Remote users complain of latency.
How should this application be redesigned to be scalable and allow for automatic failover at the lowest cost?

A. Use Amazon Route 53 failover routing with geolocation-based routin
B. Host the website on automatically scaled Amazon EC2 instances behind an Application Load Balancer with an additional Application Load Balancer and EC2 instances for the application layer in each regio
C. Use a Multi-AZ deployment with MySQL as the data layer.
D. Use Amazon Route 53 round robin routing to distribute the load evenly to several regions with health check
E. Host the website on automatically scaled Amazon ECS with AWS Fargate technology containers behind a Network Load Balancer, with an additional Network Load Balancer and Fargate containers for the application layer in each regio
F. Use Amazon Aurora replicas for the data layer.
G. Use Amazon Route 53 latency-based routing to route to the nearest region with health check
H. Host the website in Amazon S3 in each region and use Amazon API Gateway with AWS Lambda for the application laye
I. Use Amazon DynamoDB global tables as the data layer with Amazon DynamoDB Accelerator (DAX) for caching.
J. Use Amazon Route 53 geolocation-based routin
K. Host the website on automatically scaled AWS Fargate containers behind a Network Load Balancer with an additional Network Load Balancer and Fargate containers for the application layer in each regio
L. Use Amazon Aurora Multi-Master for Aurora MySQL as the data layer.

Answer: C

Explanation:
https://aws.amazon.com/getting-started/projects/build-serverless-web-app-lambda-apigateway-s3-dynamodb-co

NEW QUESTION 68
A company wants to manage the costs associated with a group of 20 applications that are critical, by migrating to AWS. The applications are a mix of Java and Node.js spread across different instance clusters. The company wants to minimize costs while standardizing by using a single deployment methodology. Most of the applications are part of month-end processing routines with a small number of concurrent users, but they are occasionally run at other times. Average application memory consumption is less than 1 GB, though some applications use as much as 2.5 GB of memory during peak processing. The most important application in the group is a billing report written in Java that accesses multiple data sources and often for several hours.
Which is the MOST cost-effective solution?

A. Deploy a separate AWS Lambda function for each applicatio
B. Use AWS CloudTrail logs and Amazon CloudWatch alarms to verify completion of critical jobs.
C. Deploy Amazon ECS containers on Amazon EC2 with Auto Scaling configured for memory utilization of 75%. Deploy an ECS task for each application being migrated with ECS task scalin
D. Monitor services and hosts by using Amazon CloudWatch.
E. Deploy AWS Elastic Beanstalk for each application with Auto Scaling to ensure that all requests have sufficient resource
F. Monitor each AWS Elastic Beanstalk deployment with using CloudWatch alarms.
G. Deploy a new amazon EC2 instance cluster that co-hosts all applications by using EC2 Auto Scaling and Application Load Balancer
H. Scale cluster size based on a custom metric set on instance memory utilizatio
I. Purchase 3-year Reserved instance reservations equal to the GroupMaxSize parameter of the Auto Scaling group.

Answer: C

NEW QUESTION 72
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 create
B. Train users to launch the template from the CloudFormation console.
C. Create an AWS Service Catalog product form the environment templat
D. Add a launch constraint to the product with the existing rol
E. Give users in the QA department permission to use AWS Service Catalog APIs onl
F. Train users to launch the templates form 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 create
H. Train users to launch the template form the CloudFormation console.
I. Create an AWS Elastic Beanstalk application from the environment templat
J. Give users in the QA department permission to use Elastic Beanstalk permissions onl
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 73
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 sharin
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 sharin
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 76
An online e-commerce business is running a workload on AWS. The application architecture includes a web tier, an application tier for business logic, and a database tier for user and transactional data management. The database server has a 100 GB memory requirement. The business requires cost-efficient disaster recovery for the application with an RTO of 5 minutes and an RPO of 1 hour. The business also has a regulatory for out-of region disaster recovery with a minimum distance between the primary and alternate sites of 250 miles.
Which of the following options can the Solutions Architect design to create a comprehensive solution for this customer that meets the disaster recovery requirements?

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

Answer: C

NEW QUESTION 79
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 82
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 templat

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 85
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 instanc
B. Create a script that will check and repair the database upon reboo
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 tw
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 aminimum instance count of on
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 loa
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 88
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 89
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 gatewa
C. Use the catalog of faces to build a collection in Amazon Rekognitio
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 gatewa
G. Use the catalog of faces to build a collection in Amazon Rekognitio
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 Rekognitio
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 solutio
N. Configure the stream to store the videos in Amazon S3.
O. Set up an Amazon EC2 instance that runs the OpenCV librarie
P. Copy the videos, images, and face catalog from the on-premises library into an Amazon EBS volume mounted on this EC2 instanc
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 91
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 dat
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 dat
F. The policy should allow the accounts in the partnership read access to the bucke
G. Enable Requester Pays on the bucke
H. Have the organizations use their AWS credentials whenaccessing 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 bucke
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 dat
O. Enable Requester Pays on the bucke
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 94
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 process. 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 template 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 96
A company has a High Performance Computing (HPC) cluster in its on-premises data center which runs thousands of jobs in parallel for one week every month, processing petabytes of images. The images are stored on a network file server, which is replicated to a disaster recovery site. The on-premises data center has reached capacity and has started to spread the jobs out over the course of month in order to better utilize the cluster, causing a delay in the job completion.
The company has asked its Solutions Architect to design a cost-effective solution on AWS to scale beyond the current capacity of 5,000 cores and 10 petabytes of data. The solution must require the least amount of management overhead and maintain the current level of durability.
Which solution will meet the company's requirements?

A. Create a container in the Amazon Elastic Container Registry with the executable file for the jo
B. Use Amazon ECS with Spot Fleet in Auto Scaling group
C. Store the raw data in Amazon EBS SC1 volumes and write the output to Amazon S3.
D. Create an Amazon EMR cluster with a combination of On Demand and Reserved Instance Task Nodes that will use Spark to pull data from Amazon S3. Use Amazon DynamoDB to maintain a list of jobs that need to be processed by the Amazon EMR cluster.
E. Store the raw data in Amazon S3, and use AWS Batch with Managed Compute Environments to create Spot Fleet
F. Submit jobs to AWS Batch Job Queues to pull down objects from Amazon S3 onto Amazon EBS volumes for temporary storage to be processed, and then write the results back to Amazon S3.
G. Submit the list of jobs to be processed to an Amazon SQS to queue the jobs that need to be processed.Create a diversified cluster of Amazon EC2 worker instances using Spot Fleet that will automatically scale based on the queue dept
H. Use Amazon EFS to store all the data sharing it across all instances in the cluster.

Answer: B

NEW QUESTION 99
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 DM
B. When the database is available in Amazon 3, 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 AW
D. Cutover to Amazon RDS after the data is synchronized.
E. Order an AWS Snowball device and copy a database dump to the devic
F. After the data has been copied to Amazon S3, import it to the Amazon RDS instanc
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 103
A company's application is increasingly popular and experiencing latency because of high volume reads on the database server.
The service has the following properties:

≫ A highly available REST API hosted in one region using Application Load Balancer (ALB) with auto scaling.

≫ A MySQL database hosted on an Amazon EC2 instance in a single Availability Zone.
The company wants to reduce latency, increase in-region database read performance, and have multi-region disaster recovery capabilities that can perform a live recovery automatically without any data or performance loss (HA/DR).
Which deployment strategy will meet these requirements?

A. Use AWS CloudFormation StackSets to deploy the API layer in two region
B. Migrate the database to an Amazon Aurora with MySQL database cluster with multiple read replicas in one region and a read replica in a different region than the source database cluste
C. Use Amazon Route 53 health checks to trigger a DNS failover to the standby region if the health checks to the primary load balancer fai
D. In the event of Route 53 failover, promote the cross-region database replica to be the master and build out new read replicas in the standby region.
E. Use Amazon ElastiCache for Redis Multi-AZ with an automatic failover to cache the database readquerie
F. Use AWS OpsWorks to deploy the API layer, cache layer, and existing database layer in two region
G. In the event of failure, use Amazon Route 53 health checks on the database to trigger a DNS failover to the standby region if the health checks in the primary region fai
H. Back up the MySQL database frequently, and in the event of a failure in an active region, copy the backup to the standby region and restore the standby database.
I. Use AWS CloudFormation StackSets to deploy the API layer in two region
J. Add the database to an Auto Scaling grou
K. Add a read replica to the database in the second regio
L. Use Amazon Route 53 health checks on the database to trigger a DNS failover to the standby region if the health checks in the primary region fai
M. Promote the cross-region database replica to be the master and build out new read replicas in the standby region.
N. Use Amazon ElastiCache for Redis Multi-AZ with an automatic failover to cache the database read querie
O. Use AWS OpsWorks to deploy the API layer, cache layer, and existing database layer in two region
P. Use Amazon Route 53 health checks on the ALB to trigger a DNS failover to the standby region if the health checks in the primary region fai
Q. Back up the MySQL database frequently, and in the event of a failure in an active region, copy the backup to the standby region and restore the standby database.

Answer: A

NEW QUESTION 105
A company is using AWS to run an internet-facing production application written in Node.js. The Development team is responsible for pushing new versions of their software directly to production. The application software is updated multiple times a day. The team needs guidance from a Solutions Architect to help them deploy the software to the production fleet quickly and with the least amount of disruption to the service.
Which option meets these requirements?

A. Prepackage the software into an AMI and then use Auto Scaling to deploy the production flee
B. For software changes, update the AMI and allow Auto Scaling to automatically push the new AMI to production.
C. Use AWS CodeDeploy to push the prepackaged AMI to productio
D. For software changes, reconfigure CodeDeploy with new AMI identification to push the new AMI to the production fleet.
E. Use AWS Elastic Beanstalk to host the production applicatio
F. For software changes, upload the new application version to Elastic Beanstalk to push this to the production fleet using a blue/green deployment method.
G. Deploy the base AMI through Auto Scaling and bootstrap the software using user dat
H. For software changes, SSH to each of the instances and replace the software with the new version.

Answer: C

NEW QUESTION 108
A company has a legacy application running on servers on premises. To increase the application's reliability, the company wants to gain actionable insights using application logs. A Solutions Architect has been given following requirements for the solution:

≫ Aggregate logs using AWS.

≫ Automate log analysis for errors.

≫ Notify the Operations team when errors go beyond a specified threshold. What solution meets the requirements?

A. Install Amazon Kinesis Agent on servers, send logs to Amazon Kinesis Data Streams and use Amazon Kinesis Data Analytics to identify errors, create an Amazon CloudWatch alarm to notify the Operations team of errors

B. Install an AWS X-Ray agent on servers, send logs to AWS Lambda and analyze them to identify errors, use Amazon CloudWatch Events to notify the Operations team of errors.

C. Install Logstash on servers, send logs to Amazon S3 and use Amazon Athena to identify errors, use sendmail to notify the Operations team of errors.

D. Install the Amazon CloudWatch agent on servers, send logs to Amazon CloudWatch Logs and use metric filters to identify errors, create a CloudWatch alarm to notify the Operations team of errors.

Answer: A

Explanation:
https://docs.aws.amazon.com/kinesis-agent-windows/latest/userguide/what-is-kinesis-agent-windows.html https://medium.com/@khandelwal12nidhi/build-log-analytic-solution-on-aws-cc62a70057b2

NEW QUESTION 112
An enterprise company is using a multi-account AWS strategy There are separate accounts tor development staging and production workloads To control costs and improve governance the following requirements have been defined:
• The company must be able to calculate the AWS costs tor each project
• The company must be able to calculate the AWS costs tor each environment development staging and production
• Commonly deployed IT services must be centrally managed
• Business units can deploy pre-approved IT services only
• Usage of AWS resources in the development account must be limited
Which combination of actions should be taken to meet these requirements? (Select THREE)

A. Apply environment, cost center, and application name tags to all taggable resources
B. Configure custom budgets and define thresholds using Cost Explorer
C. Configure AWS Trusted Advisor to obtain weekly emails with cost-saving estimates
D. Create a portfolio for each business unit and add products to the portfolios using AWS CloudFormation in AWS Service Catalog
E. Configure a billing alarm in Amazon CloudWatch.
F. Configure SCPs in AWS Organizations to allow services available using AWS

Answer: CEF

NEW QUESTION 114
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 117
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 failure 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 Redshif
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 Forehose delivery strea
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-tim
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-tim
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 strea
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 119
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 compan
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 compan
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 invoic
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 acces
H. Create policies that allow the users access to all resources in that accoun
I. Attach the policies to the IAM user.
J. Create a federated identity store against the company's Active Director
K. Create IAM roles with appropriate permissions and set the trust relationships with AWS and the identity stor
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 compan
N. For billing purposes, use a tagging solution that uses a tag to identify the company that creates each resource.

Answer: AD

NEW QUESTION 121
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 123
The company Security team requires that all data uploaded into an Amazon S3 bucket must be encrypted. The encryption keys must be highly available and the company must be able to control access on a per-user basis, with different users having access to different encryption keys.
Which of the following architectures will meet these requirements? (Choose two.)

A. Use Amazon S3 server-side encryption with Amazon S3-managed key
B. Allow Amazon S3 to generate an AWS/S3 master key, and use IAM to control access to the data keys that are generated.
C. Use Amazon S3 server-side encryption with AWS KMS-managed keys, create multiple customer master keys, and use key policies to control access to them.
D. Use Amazon S3 server-side encryption with customer-managed keys, and use AWS CloudHSM to manage the key
E. Use CloudHSM client software to control access to the keys that are generated.
F. Use Amazon S3 server-side encryption with customer-managed keys, and use two AWS CloudHSM instances configured in high-availability mode to manage the key
G. Use the Cloud HSM client software to control access to the keys that are generated.
H. Use Amazon S3 server-side encryption with customer-managed keys, and use two AWS CloudHSM instances configured in high-availability mode to manage the key
I. Use IAM to control access to the keys that are generated in CloudHSM.

Answer: BD

Explanation:
http://websecuritypatterns.com/blogs/2018/03/01/encryption-and-key-management-in-aws-kms-vs-cloudhsm-mys/

NEW QUESTION 128
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 instance
B. Ensure that all Windows EC2 instances are assigned this ta
C. Associate the AWS-DefaultPatchBaseline to the Windows servers patch grou
D. Define an AWS Systems Manager maintenance window, conduct patching within it, and associate it with the Windows Servers patch grou
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 a value of Windows Servers to all existing EC2 instance
H. Ensure that all Windows EC2 instances are assigned this ta
I. Associate the AWS-WindowsPatchBaseline document as a task associated with the Windows Servers patch grou

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 comman

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 Server2 to all existing EC2 instance

M. Ensure that all Windows EC2 instances are assigned this ta

N. Associate theAWS-DefaultPatchBaseline with both Windows Servers patch group

O. Define two non-overlappingAWS Systems Manager maintenance windows, conduct patching within them, and associate each with a different patch grou

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 Server2 to all existing EC2 instance

S. Ensure that all Windows EC2 instances are assigned this ta

T. Associate theAWS-WindowsPatchBaseline with both Windows Servers patch group

. Define two non-overlappingAWS Systems Manager maintenance windows, conduct patching within them, and associate each with a different patch grou

. Assign the AWS-RunWindowsPatchBaseline document as a task within each maintenance windo

. Create an AWS Systems Manager State Manager document to define commands to be executed during patch execution.

Answer: C

NEW QUESTION 133
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 configuratio

B. Ensure that all EC2 instances are purchased as reserved instance

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 EC

E. Migrate to an Amazon Aurora MySQL Multi-AZ cluste

F. Implement storage checks for Aurora MySQL storage utilization and an AWS Lambda function to grow the Aurora MySQL storage, as necessar

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 instance

I. Ensure that the reserved instances are purchased for fixed capacity and that Auto Scaling instances run on deman

J. Migrate to an Amazon Aurora MySQLMulti-AZ cluste

K. Ensure that Multi-AZ architectures are implemented.

L. Ensure that EC2 instances are right-sized and behind an Elastic Load Balance

M. Implement Auto Scaling with EC2 instance

N. Ensure that Reserved instances are purchased for fixed capacity and that Auto Scaling instances run on deman

O. Migrate to an Amazon Aurora MySQL Multi-AZ cluste

P. Implement storage checks for Aurora MySQL storage utilization and an AWS Lambda function to grow Aurora MySQL storage, as necessar

Q. Ensure Multi-AZ architectures are implemented.

Answer: C

NEW QUESTION 138
A Solutions Architect must establish a patching plan for a large mixed fleet of Windows and Linux servers. The patching plan must be implemented securely, be audit ready, and comply with the company's business requirements.
Which option will meet these requirements with MINIMAL effort?

A. Install and use an OS-native patching service to manage the update frequency and release approval for all instance

B. Use AWS Config to verify the OS state on each instance and report on any patch compliance issues.

C. Use AWS Systems Manager on all instances to manage patchin

D. Test patches outside of production and then deploy during a maintenance window with the appropriate approval.

E. Use AWS OpsWorks for Chef Automate to run a set of scripts that will iterate through all instances of a given typ

F. Issue the appropriate OS command to get and install updates on each instance, including any required restarts during the maintenance window.

G. Migrate all applications to AWS OpsWorks and use OpsWorks automatic patching support to keep the OS up-to-date following the initial installatio

H. Use AWS Config to provide audit and compliance reporting.

Answer: B

Explanation:
Only Systems Manager can patch both OS effectively on AWS and on premise.

NEW QUESTION 142
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 serve

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 serve

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 143

A company's main intranet page has experienced degraded response times as its user base has increased although there are no reports of users seeing error pages. The application uses Amazon DynamoDB in read-only mode.
Amazon DynamoDB latency metrics for successful requests have been in a steady state even during times when users have reported degradation The Development team has correlated the issue to ProvisionedThrough put Exceeded exceptions in the application logs when doing Scan and read operations The team also identified an access pattern of steady spikes of read activity on a distributed set of individual data items
The Chief Technology Officer wants to improve the user experience
Which solutions will meet these requirements with the LEAST amount of changes to the application? (Select TWO)

A. Change the data model of the DynamoDB tables to ensure that all Scan and read operations meet DynamoDB best practices of uniform data access, reaching the full request throughput provisioned for the DynamoDB tables
B. Enable DynamoDB auto scaling to manage the throughput capacity as table traffic increases Set the upper and lower limits to control costs and set a target utilization given the peak usage and how quickly the traffic changes.
C. Provision Amazon ElastiCache for Redis with cluster mode enabled The cluster should be provisioned with enough shards to spread the application load and provision at least one read replica node for each shard
D. Implement the DynamoDB Accelerator (DAX) client and provision a DAX cluster with the appropriate node types to sustain the application loa
E. Tune the item and query cache configuration for an optimal user experience
F. Remove error retries and exponential backoffs in the application code to handle throttling errors

Answer: AE

NEW QUESTION 146
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 migratio
B. Use AWS DMS to analyse the current schema and provide a recommendation for the optimal database engin
C. Then, use AWS DMS to migrate to the recommended enginee
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 migratio
F. Use AWS DMS to begin moving data from the on-premises database to AW
G. After the initial copy, continue to use AWS DMS to keep the databases in sync until cutting over to the new databas
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 RD
J. Then, use AWS DMS to migrate to the platfor
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 AW
M. After the initial copy, continue to use AWS DMS to keep the databases in sync until cutting over to the new databas
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 147
A Solution Architect is designing a deployment strategy for an application tier and gas the following requirements.
* The application code will need a 500 HB static dataset to be present before application startup.
* The application tier be able to scale Up and down based on demand with as little startup time as possible.
* The development team should be able to update the code multiple times each day.
* Critical operating system (OS) patches must be installed within 48 hours of being released. Which deployment strategy meets these requirements?

A. Use AWS Manager to create a new AMI with the updated OS patches . Update the Auto Scaling group to use the patches AMI and replace existing unpatche
B. Use AWS CodeDeploy to push the application code to the instance
C. Store the static data in Amazon EFS.
D. Use AWS System Manager to create a new AMI with upload OS patche
E. Update the Auto Scaling group to use the patches AMI and replace existing unpatches and the application code as a batch job every nigh
F. Store the static data in Amazon EFS.
G. Use an Amazon provided AMI for the OS Configure an Auto Scaling group set to a static instance coun
H. Configure an Amazon EC2 data script to download the data from Amazon S3 install OS patches with AWS system Manager when they are release
I. Use Codedeploy to push the application code to the instances.
J. Use an Amazon provided AMI for the OS Configure an Auto Scaling group Configure an Amazon EC2 user data script to download the data from Amazon S3. Replace existing instances after eachAmazon-provided AMI releas
K. Use AWS CodeDeploy to push the application code to the instances.

Answer: C

NEW QUESTION 149
A Solutions Architect is working with a company that operates a standard three-tier web application in AWS. The web and application tiers run on Amazon EC2 and the database tier runs on Amazon RDS. The company is redesigning the web and application tiers to use Amazon API Gateway and AWS Lambda, and the company intends to deploy the new application within 6 months. The IT Manager has asked the Solutions Architect to reduce costs in the interim.
Which solution will be MOST cost effective while maintaining reliability?

A. Use Spot Instances for the web tier, On-Demand Instances for the application tier, and Reserved Instances for the database tier.
B. Use On-Demand Instances for the web and application tiers, and Reserved Instances for the database tier.
C. Use Spot Instances for the web and application tiers, and Reserved Instances for the database tier.
D. Use Reserved Instances for the web, application, and database tiers.

Answer: B

NEW QUESTION 152

A retail company is running an application that stores invoice files in Amazon S3 bucket and metadata about the files in an Amazon. The S3 bucket and DynamoDB table are in us-east-1. The company wants to protect itself from data corruption and loss of connectivity to either Region.
Which option meets these requirements?

A. Create a DynamoDB global table to replicate data between us-east-1 and eu-west-1. Enable continuous backup on the DynamoDB table in us-east-1. Enable versioning on the S3 bucket.
B. Create an AWS Lambda function triggered by Amazon CloudWatch Events to make regular backups of the DynamoDB tabl
C. Set up S3 cross-region replication from us-east-1 to eu-west-1. Set up MFA deleteon the S3 bucket in us-east-1.
D. Create a DynamoDB global table to replicate data between us-east-1 and eu-west-1. Enable versioning on the S3 bucke
E. Implement strict ACLs on the S3 bucket.
F. Create a DynamoDB global table to replicate data between us-east-1 and eu-west-1. Enable continuous backup on the DynamoDB table in us-east-1. Set up S3 cross-region replication from us-east-1 toeu-west-1.

Answer: A

Explanation:
https://aws.amazon.com/blogs/aws/new-cross-region-replication-for-amazon-s3/

NEW QUESTION 154
A Solutions Architect is designing a multi-account structure that has 10 existing accounts. The design must meet the following requirements:

> Consolidate all accounts into one organization.

> Allow full access to the Amazon EC2 service from the master account and the secondary accounts.

> Minimize the effort required to add additional secondary accounts.
Which combination of steps should be included in the solution? (Choose two.)

A. Create an organization from the master accoun
B. Send invitations to the secondary accounts from the master accoun
C. Accept the invitations and create an OU.
D. Create an organization from the master accoun
E. Send a join request to the master account from each secondary accoun
F. Accept the requests and create an OU.
G. Create a VPC peering connection between the master account and the secondary account
H. Accept the request for the VPC peering connection.
I. Create a service control policy (SCP) that enables full EC2 access, and attach the policy to the OU.
J. Create a full EC2 access policy and map the policy to a role in each accoun
K. Trust every other account to assume the role.

Answer: AD

Explanation:
There is a concept of Permission Boundary vs Actual IAM Policies That is, we have a concept of "Allow" vs "Grant". In terms of boundaries, we have the following three boundaries: 1. SCP 2. User/Role boundaries 3. Session boundaries (ex. AssumeRole ...) In terms of actual permission granting, we have the following: 1. Identity Policies 2. Resource Policies

NEW QUESTION 157
A company is migrating its on-premises build artifact server to an AWS solution. The current system consists of an Apache HTTP server that serves artifacts to clients on the local network, restricted by the perimeter firewall. The artifact consumers are largely build automation scripts that download artifacts via anonymous HTTP, which the company will be unable to modify within its migration timetable.
The company decides to move the solution to Amazon S3 static website hosting. The artifact consumers will be migrated to Amazon EC2 instances located within both public and private subnets in a virtual private cloud (VPC).
Which solution will permit the artifact consumers to download artifacts without modifying the existing automation scripts?

A. Create a NAT gateway within a public subnet of the VP
B. Add a default route pointing to the NAT gateway into the route table associated with the subnets containing consumer
C. Configure the bucket policy to allow the s3:ListBucket and s3:GetObject actions using the condition IpAddress and the condition key aws:SourceIp matching the elastic IP address if the NAT gateway.
D. Create a VPC endpoint and add it to the route table associated with subnets containing consumers.Configure the bucket policy to allow s3:ListBucket and s3:GetObject actions using the condition StringEquals and the condition key aws:sourceVpce matching the identification of the VPC endpoint.
E. Create an IAM role and instance profile for Amazon EC2 and attach it to the instances that consume build artifact
F. Configure the bucket policy to allow the s3:ListBucket and s3:GetObjects actions for theprincipal matching the IAM role created.
G. Create a VPC endpoint and add it to the route table associated with subnets containing consumers.Configure the bucket policy to allow s3:ListBucket and s3:GetObject actions using the condition IpAddress and the condition key aws:SourceIp matching the VPC CIDR block.

Answer: B

NEW QUESTION 162
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 regio
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 regio
D. Use a lifecycle policy to move this data to Amazon Glacier in the production region immediatel
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 regio
G. Set up cross-region replication of this data to Amazon Glacier in the disaster recovery regio
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 regio
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 tor the production accounts
C. Create a SCP containing a Deny Effect for cloudfront". Iam:*, 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'. Iam * 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 168
A company is running a high-user-volume media-sharing application on premises It currently hosts about 400 TB of data with millions of video files The company is migrating this application to AWS to improve reliability and reduce costs
The Solutions Architecture team plans to store the videos in an Amazon S3 bucket and use Amazon
CloudFront to distribute videos to users. The company needs to migrate this application to AWS within 10 days with the least amount of downtime possible. The company currently has 1 Gbps connectivity to the internet with 30 percent free capacity
Which of the following solutions would enable the company to migrate the workload to AWS and meet an of the requirements?

A. Use a multipart upload in Amazon S3 client at to parallel-upload the data to the Amazon S3 bucket over the internet Use the throttling feature to ensure that the Amazon S3 client does not use more than 30 percent of available internet capacity
B. Request an AWS Snowmobile with 1 PB capacity to be delivered to the data center Load the data into Snowmobile and send it back to have AWS download that data to the Amazon S3 bucket Sync the new data that was generated white migration was in flight
C. Use an Amazon S3 client to transfer data from the data center to the Amazon S3 bucket over the internet Use the throttling feature to ensure the Amazon S3 client does not use more than 30 percent of available internet capacity
D. Request multiple AWS Snowball devices to be delivered to the data center Load the data concurrently into these devices and send it back Have AWS download that data to the Amazon S3 bucket Sync the new data that was generated while migration was in flight.

Answer: D

Explanation:
https://www.edureka.co/blog/aws-snowball-and-snowmobile-tutorial/

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 175
A company has asked a Solutions Architect to design a secure content management solution that can be accessed by API calls by external customer applications. The company requires that a customer administrator must be able to submit an API call and roll back changes to existing files sent to the content management solution, as needed.
What is the MOST secure deployment design that meets all solution requirements?

A. Use Amazon S3 for object storage with versioning and bucket access logging enabled, and an IAM role and access policy for each customer applicatio
B. Encrypt objects using SSE-KM
C. Develop the content management application to use a separate AWS KMS key for each customer.
D. Use Amazon WorkDocs for object storag
E. Leverage WorkDocs encryption, user access management, and version contro
F. Use AWS CloudTrail to log all SDK actions and create reports of hourly access by using the Amazon CloudWatch dashboar
G. Enable a revert function in the SDK based on a static Amazon S3 webpage that shows the output of the CloudWatch dashboard.
H. Use Amazon EFS for object storage, using encryption at rest for the Amazon EFS volume and a customer managed key stored in AWS KM
I. Use IAM roles and Amazon EFS access policies to specify separate encryption keys for each customer applicatio
J. Deploy the content management application to store all new versions as new files in Amazon EFS and use a control API to revert a specific file to a previous version.
K. Use Amazon S3 for object storage with versioning and enable S3 bucket access loggin
L. Use an IAM role and access policy for each customer applicatio
M. Encrypt objects using client-side encryption, and distribute an encryption key to all customers when accessing the content management application.

Answer: A

NEW QUESTION 179
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 184
A company has deployed an application to multiple environments in AWS, including production and testing. The company has separate accounts for production and testing, and users are allowed to create additional application users for team members or services, as needed. The Security team has asked the Operations team for better isolation between production and testing with centralized controls on security credentials and improved management of permissions between environments.
Which of the following options would MOST securely accomplish this goal?

A. Create a new AWS account to hold user and service accounts, such as an identity accoun
B. Create users and groups in the identity accoun
C. Create roles with appropriate permissions in the production and testing account
D. Add the identity account to the trust policies for the roles.
E. Modify permissions in the production and testing accounts to limit creating new IAM users to members of the Operations tea
F. Set a strong IAM password policy on each accoun
G. Create new IAM users and groups in each account to limit developer access to just the services required to complete their job function.
H. Create a script that runs on each account that checks user accounts for adherence to a security policy.Disable any user or service accounts that do not comply.
I. Create all user accounts in the production accoun
J. Create roles for access in the production account and testing account.
K. Grant cross-account access from the production account to the testing account.

Answer: A

Explanation:
https://aws.amazon.com/blogs/security/how-to-centralize-and-automate-iam-policy-creation-in-sandbox-develop

NEW QUESTION 189
A public retail web application uses an Application Load Balancer (ALB) in front of Amazon EC2 instances running across multiple Availability Zones (AZs) in a Region backed by an Amazon RDS MySQL Multi-AZ deployment. Target group health checks are configured to use HTTP and pointed at the product catalog page. Auto Scaling is configured to maintain the web fleet size based on the ALB health check.
Recently, the application experienced an outage. Auto Scaling continuously replaced the instances during the outage. A subsequent investigation determined that the web server metrics were within the normal range, but the database tier was experiencing high load, resulting in severely elevated query response times.

Which of the following changes together would remediate these issues while improving monitoring capabilities for the availability and functionality of the entire application stack for future growth? (Select TWO.)

A. Configure read replicas for Amazon RDS MySQL and use the single reader endpoint in the web application to reduce the load on the backend database tier.
B. Configure the target group health check to point at a simple HTML page instead of a product catalog page and the Amazon Route 53 health check against the product page to evaluate full application functionalit
C. Configure Amazon CloudWatch alarms to notify administrators when the site fails.
D. Configure the target group health check to use a TCP check of the Amazon EC2 web server and the Amazon Route 53 health check against the product page to evaluate full application functionalit
E. Configure Amazon CloudWatch alarms to notify administrators when the site fails.
F. Configure an Amazon CloudWatch alarm for Amazon RDS with an action to recover a high-load, impaired RDS instance in the database tier.
G. Configure an Amazon ElastiCache cluster and place it between the web application and RDS MySQL instances to reduce the load on the backend database tier.

Answer: BE

NEW QUESTION 190
......

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