SAP-C01 Dumps

AWS Certified Solutions Architect- Professional

https://www.certleader.com/SAP-C01-dumps.html

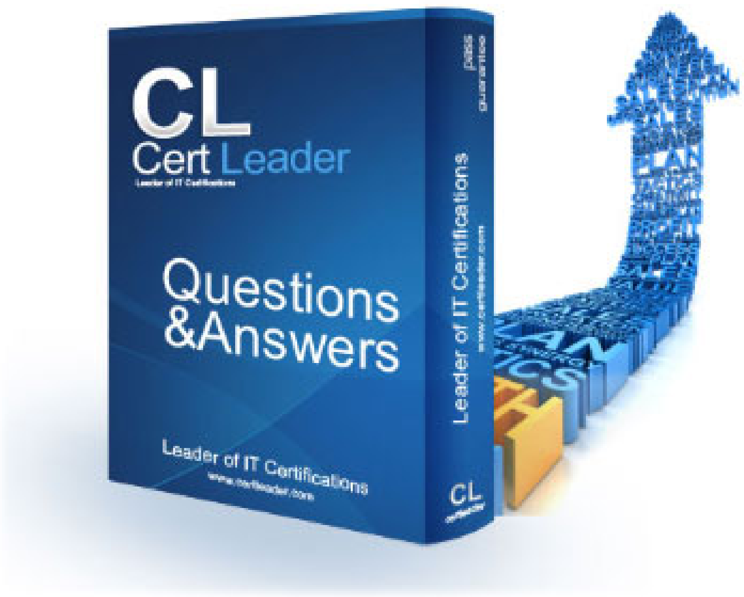

NEW QUESTION 1
A Solutions Architect is building a containerized NET Core application that will run in AWS Fargate The backend of the application requires Microsoft SQL Server with high availability All tiers of the application must be highly available The credentials used for the connection string to 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 Solutions Architect must design a highly available, stateless, REST service. The service will require multiple persistent storage layers for service object meta information and the delivery of content. Each request needs to be authenticated and securely processed. There is a requirement to keep costs as low as possible?
How can these requirements be met?

A. Use AWS Fargate to host a container that runs a self-contained REST servic
B. Set up an Amazon ECS service that is fronted by an Application Load Balancer (ALB). Use a custom authenticator to control access to the AP
C. Store request meta information in Amazon DynamoDB with Auto Scaling and static content in a secured S3 bucke
D. Make secure signed requests for Amazon S3 objects and proxy the data through the REST service interface.
E. Use AWS Fargate to host a container that runs a self-contained REST servic
F. Set up an ECS service that is fronted by a cross-zone AL
G. Use an Amazon Cognito user pool to control access to the AP
H. Store request meta information in DynamoDB with Auto Scaling and static content in a secured S3 bucke
I. Generate presigned URLs when returning references to content stored in Amazon S3.
J. Set up Amazon API Gateway and create the required API resources and method
K. Use an Amazon Cognito user pool to control access to the AP
L. Configure the methods to use AWS Lambda proxy integrations, and process each resource with a unique AWS Lambda functio
M. Store request meta information in DynamoDB with Auto Scaling and static content in a secured S3 bucke
N. Generate presigned URLs when returning references to content stored in Amazon S3.
O. Set up Amazon API Gateway and create the required API resources and method
P. Use an Amazon API Gateway custom authorizer to control access to the AP
Q. Configure the methods to use AWS Lambda custom integrations, and process each resource with a unique Lambda functio
R. Store request meta information in an Amazon ElastiCache Multi-AZ cluster and static content in a secured S3 bucke
S. Generate presigned URLs when returning references to content stored in Amazon S3.

Answer: C

NEW QUESTION 3
A retail company processes point-of-state data on application servers in its data center and writes outputs to Amazon DynamoDB table. The data center is connected to the company’s VPC with an AWS Direct Connect (DX) connection, and the application servers require a consistent network connection at speed greater than 2 Gbps.
The company decides that the DynamoDB table needs to be highly available and fault tolerant. The company policy states that the data should be available across two regions.
What changes should the company make to meet these requirements?

A. Establish a second DX connection for redundanc
B. Use DynamoDB global tables to replicate data to a second Region modify the application to fail over to the second Region.
C. Use an AWS managed VPN as a backup to D
D. Create an identical DynamoDB table in a second Regio
E. Modify the application to replicate data to both regions.
F. Establish a second DX connection for redundanc
G. Create an identical DynamoDB table in a second Regio
H. Enable DynamoDB auto scaling to manage throughput capacit
I. Modify the application to write to the second Region.
J. Use AWS managed VPN as a backup to D
K. Create an identical DynamoDB table in a second Region.Enable DynamoDB streams to capture changes to the tabl
L. Use AWS Lambda to replicate changes to the second Region.

Answer: A

NEW QUESTION 4
A large company has many business units. Each business unit has multiple AWS accounts for different purposes. The CIO of the company sees that each business unit has data that would be useful to share with other parts of the company. In total, there are about 10 PB of data that needs to be shared with users in 1,000 AWS accounts. The data is proprietary, so some of it should only be available to users with specific job types. Some of the data is used for throughput of intensive workloads, such as simulations. The number of accounts changes frequently because of new initiatives, acquisitions, and divestitures.
A Solutions Architect has been asked to design a system that will allow for sharing data for use in AWS with all of the employees in the company.
Which approach will allow for secure data sharing in 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 5
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 6
While debugging a backend application for an loT system that supports globally distributed devices a Solutions Architect notices that stale data is occasionally being sent to user devices. Devices often share data, and stale data does not cause issues in most cases However device operations are disrupted when a device reads the stale data after an update
The global system has multiple identical application stacks deployed In different AWS Regions If a user device travels out of its home geographic region it will always connect to the geographically closest AWS Region to write or read data The same data is available in all supported AWS Regions using an Amazon DynamoDB global table
What change should be made to avoid causing disruptions in device operations'?

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

Answer: B

NEW QUESTION 7
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 8
A company runs an application on a fleet of Amazon EC2 instances The application requires low latency and random access to 100 GB of data The application must be able to access the data at up to 3.000 IOPS A Development team has configured the EC2 launch template to provision a 100-GB Provisioned IOPS (PIOPS) Amazon EBS volume with 3 000 IOPS provisioned A Solutions Architect is tasked with lowering costs without impacting performance and durability Which action should be taken?

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

Answer: A

NEW QUESTION 9
As a part of building large applications in the AWS Cloud, the Solutions Architect is required to implement the perimeter security protection. Applications running on AWS have the following endpoints:

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

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

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

How should the Solutions Architect design the solution?

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

Answer: C

NEW QUESTION 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 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 13
A bank is designing an online customer service portal where customers can chat with customer service agents. The portal is required to maintain a 15-minute RPO or RTO in case of a regional disaster. Banking regulations require that all customer service chat transcripts must be preserved on durable storage for at least 7 years, chat conversations must be encrypted in-flight, and transcripts must be encrypted at rest. The Data Lost Prevention team requires that data at rest must be encrypted using a key that the team controls, rotates, and revokes.
Which design meets these requirements?

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

Answer: B

NEW QUESTION 14
A company had a tight deadline to migrate its on-premises environment to AWS. It moved over Microsoft SQL Servers and Microsoft Windows Servers using the virtual machine import/export service and rebuild other applications native to the cloud. The team created both Amazon EC2 databases and used Amazon RDS. Each team in the company was responsible for migrating their applications, and they have created individual accounts for isolation of resources. The company did not have much time to consider costs, but now it would like suggestions on reducing its AWS spend.
Which steps should a Solutions Architect take to reduce costs?

A. Enable AWS Business Support and review AWS Trusted Advisor's cost check
B. Create Amazon EC2 Auto Scaling groups for applications that experience fluctuating deman
C. Save AWS Simple Monthly Calculator reports in Amazon S3 for trend analysi
D. Create a master account under Organizations and have teams join for consolidating billing.
E. Enable Cost Explorer and AWS Business Support Reserve Amazon EC2 and Amazon RDS DB instance
F. Use Amazon CloudWatch and AWS Trusted Advisor for monitoring and to receive cost-savings suggestion
G. Create a master account under Organizations and have teams join for consolidated billing.
H. Create an AWS Lambda function that changes the instance size based on Amazon CloudWatch alarms.Reserve instances based on AWS Simple Monthly Calculator suggestion
I. Have an AWSWell-Architected framework review and apply recommendation
J. Create a master account under Organizations and have teams join for consolidated billing.
K. Create a budget and monitor for costs exceeding the budge
L. Create Amazon EC2 Auto Scaling groups for applications that experience fluctuating deman
M. Create an AWS Lambda function that changes instance sizes based on Amazon CloudWatch alarm
N. Have each team upload their bill to an Amazon S3 bucket for analysis of team spendin
O. Use Spot instances on nightly batch processing jobs.

Answer: B

Explanation:
Import/Export supports importing and exporting data into and out of Amazon S3 buckets. For significant data sets, AWS Import/Export is often faster than Internet transfer and more cost effective than upgrading your connectivity.

NEW QUESTION 18
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 22
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 24
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 29
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 32
A retail company has a custom NET web application running on AWS that uses Microsoft SQL Server for the database The application servers maintain a user's session locally.
Which combination of architecture changes are needed ensure all tiers of the solution are highly available? (Select THREE.)

A. Refactor the application to store the user's session in Amazon ElastiCache Use Application Load Balancers to distribute the load between application instances
B. Set up the database to generate hourly snapshots using Amazon EBS Configure an Amazon CloudWatch Events rule to launch a new database instance if the primary one fails
C. Migrate the database to Amazon RDS tor SQL Server Configure the RDS instance to use a Multi-AZ deployment
D. Move the NET content to an Amazon S3 bucket Configure the bucket for static website hosting
E. Put the application instances in an Auto Scaling group Configure the Auto Scaling group to create new instances if an instance becomes unhealthy
F. Deploy Amazon CloudFront in front of the application tier Configure CloudFront to serve content from healthy application instances only

Answer: BDE

NEW QUESTION 35
A company runs an ordering system on AWS using Amazon SQS and AWS Lambda, with each order received as a JSON message. recently the company had a marketing event that led to a tenfold increase in orders. With this increase, the following undesired behaviors started in the ordering system:
- Lambda failures while processing orders lead to queue backlogs.
- The same orders have been processed multiple times.

A solutions Architect has been asked to solve the existing issues with the ordering system and add the following resiliency features:
- Retain problematic orders for analysis.
- Send notification if errors go beyond a threshold value. How should the Solutions Architect meet these requirements?

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

Answer: D

NEW QUESTION 40
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 41
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 43
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 48
A large company experienced a drastic increase in its monthly AWS spend. This is after Developers accidentally launched Amazon EC2 instances in unexpected regions. The company has established practices around least privileges for Developers and controls access to on-premises resources using Active Directory groups. The company now wants to control costs by restricting the level of access that Developers have to the AWS Management Console without impacting their productivity. The company would also like to allow Developers to launch Amazon EC2 in only one region, without limiting access to other services in any region.
How can this company achieve these new security requirements while minimizing the administrative burden on the Operations team?

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

Answer: D

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

NEW QUESTION 52
An enterprise runs 103 line-of-business applications on virtual machines in an on-premises data center. Many of the applications are simple PHP, Java, or Ruby web applications, are no longer actively developed, and serve little traffic.
Which approach should be used to migrate these applications to AWS with the LOWEST infrastructure costs ?

A. Deploy the applications to single-instance AWS Elastic Beanstalk environments without a load balancer.
B. Use AWS SMS to create AMIs for each virtual machine and run them in Amazon EC2.
C. Convert each application to a Docker image and deploy to a small Amazon ECS cluster behind an Application Load Balancer.
D. Use VM Import/Export to create AMIs for each virtual machine and run them in single-instance AWS Elastic Beanstalk environments by configuring a custom image.

Answer: A

Explanation:
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-types.html

NEW QUESTION 54
A company has created an account for individual Development teams, resulting in a total of 200 accounts. All accounts have a single virtual private cloud (VPC) in a single region with multiple microservices running in Docker containers that need to communicate with microservices in other accounts. The Security team requirements state that these microservices must not traverse the public internet, and only certain internal services should be allowed to call other individual services. If there is any denied network traffic for a service, the Security team must be notified of any denied requests, including the source IP.
How can connectivity be established between services while meeting the security requirements?

A. Create a VPC peering connection between the VPC
B. Use security groups on the instances to allow traffic from the security group IDs that are permitted to call the microservic
C. Apply network ACLs to and allow traffic from the local VPC and peered VPCs onl
D. Within the task definition in Amazon ECS for each of the microservices, specify a log configuration by using the awslogs drive
E. Within Amazon CloudWatch Logs, create a metric filter and alarm off of the number of HTTP 403 response
F. Create an alarm when the number of messages exceeds a threshold set by the Security team.
G. Ensure that no CIDR ranges are overlapping, and attach a virtual private gateway (VGW) to each VPC.Provision an IPsec tunnel between each VGW and enable route propagation on the route tabl
H. Configure security groups on each service to allow the CIDR ranges of the VPCs on the other account
I. Enable VPC Flow Logs, and use an Amazon CloudWatch Logs subscription filter for rejected traffi
J. Create an IAM role and allow the Security team to call the AssumeRole action for each account.
K. Deploy a transit VPC by using third-party marketplace VPN appliances running on Amazon EC2, dynamically routed VPN connections between the VPN appliance, and the virtual private gateways (VGWs) attached to each VPC within the regio
L. Adjust network ACLs to allow traffic from the local VPC onl
M. Apply security groups to the microservices to allow traffic from the VPN appliances onl
N. Install the awslogs agent on each VPN appliance, and configure logs to forward to Amazon CloudWatch Logs in the security account for the Security team to access.
O. Create a Network Load Balancer (NLB) for each microservic
P. Attach the NLB to a PrivateLink endpoint service and whitelist the accounts that will be consuming this servic
Q. Create an interface endpoint in the consumer VPC and associate a security group that allows only the security group IDs of the services authorized to call the producer servic
R. On the producer services, create security groups for each microservice and allow only the CIDR range the allowed service
S. Create VPC Flow Logs on each VPC to capture rejected traffic that will be delivered to an Amazon CloudWatch Logs grou
T. Create a CloudWatch Logs subscription that streams the log data to a security account.

Answer: D

Explanation:

AWS PrivateLink provides private connectivity between VPCs, AWS services, and on-premises applications, securely on the Amazon network. AWS PrivateLink makes it easy to connect services across different accounts and VPCs to significantly simplify the network architecture. It seems like the next VPC peering.
https://aws.amazon.com/privatelink/

NEW QUESTION 58
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 60
A company runs a video processing platform. Files are uploaded by users who connect to a web server, which stores them on an Amazon EFS share. This web server is running on a single Amazon EC2 instance. A different group of instances, running in an Auto Scaling group, scans the EFS share directory structure for new files to process and generates new videos (thumbnails, different resolution, compression, etc.) according to the instructions file, which is uploaded along with the video files. A different application running on a group of instances managed by an Auto Scaling group processes the video files and then deletes them from the EFS share. The results are stored in an S3 bucket. Links to the processed video files are emailed to the customer.
The company has recently discovered that as they add more instances to the Auto Scaling Group, many files are processed twice, so image processing speed is not improved. The maximum size of these video files is 2GB.
What should the Solutions Architect do to improve reliability and reduce the redundant processing of video files?

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

Answer: A

NEW QUESTION 65
A company is moving a business-critical application onto AWS. It is a traditional three-tier web application using an Oracle database. Data must be encrypted in transit and at rest. The database hosts 12 TB of data. Network connectivity to the source Oracle database over the internal is allowed, and the company wants to reduce the operational costs by using AWS Managed Services where possible. All resources within the web and application tiers have been migrated. The database has a few tables and a simple schema using primary keys only; however, it contains many Binary Large Object (BLOB) fields. It was not possible to use the database's native replication tools because of licensing restrictions.
Which database migration solution will result in the LEAST amount of impact to the application's availability?

A. Provision an Amazon RDS for Oracle instanc
B. Host the RDS database within a virtual private cloud (VPC) subnet with internet access, and set up the RDS database as an encrypted Read Replica of the source databas
C. Use SSL to encrypt the connection between the two database
D. Monitor the replication performance by watching the RDS ReplicaLag metri
E. During the application maintenance window, shut down the on-premises database and switch over the application connection to the RDS instance when there is no more replication la
F. Promote the Read Replica into a standalone database instance.
G. Provision an Amazon EC2 instance and install the same Oracle database softwar
H. Create a backup of the source database using the supported tool
I. During the application maintenance window, restore the backup into the Oracle database running in the EC2 instanc
J. Set up an Amazon RDS for Oracle instance, and create an import job between the database hosted in AW
K. Shut down the source database and switch over the database connections to the RDS instance when the job is complete.
L. Use AWS DMS to load and replicate the dataset between the on-premises Oracle database and the replication instance hosted on AW
M. Provision an Amazon RDS for Oracle instance with Transparent Data Encryption (TDE) enabled and configure it as target for the replication instanc

N. Create a customer-managed AWS KMS master key to set it as the encryption key for the replication instance.Use AWS DMS tasks to load the data into the target RDS instanc
O. During the application maintenance window and after the load tasks reach the ongoing replication phase, switch the database connections to the new database.
P. Create a compressed full database backup on the on-premises Oracle database during an application maintenance windo
Q. While the backup is being performed, provision a 10 Gbps AWS Direct Connect connection to increase the transfer speed of the database backup files to Amazon S3, and shorten the maintenance window perio
R. Use SSL/TLS to copy the files over the Direct Connect connectio
S. When the backup files are successfully copied, start the maintenance window, and rise any of the Amazon RDS supported tools to import the data into a newly provisioned Amazon RDS for Oracle instance with encryption enable
T. Wait until the data is fully loaded and switch over the database connections to the new databas
. Delete the Direct Connect connection to cut unnecessary charges.

Answer: C

Explanation:
https://aws.amazon.com/blogs/apn/oracle-database-encryption-options-on-amazon-rds/
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Appendix.Oracle.Options.AdvSecurity.htm l (DMS in transit encryption)
https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Security.html

NEW QUESTION 68
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 70
A finance company is running its business-critical application on current-generation Linux EC2 instances. The application includes a self-managed MySQL database performing heavy I/O operations. The application is working fine to handle a moderate amount of traffic during the month. However, it slows down during the final three days of each month due to month-end reporting, even though the company is using Elastic Load Balancers and Auto Scaling within its infrastructure to meet the increased demand.
Which of the following actions would allow the database to handle the month-end load with the LEAST impact on performance?

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

Answer: B

NEW QUESTION 72
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 75
A company manages more than 200 separate internet-facing web applications. All of the applications are deployed to AWS in a single AWS Region The fully qualified domain names (FQDNs) of all of the applications are made available through HTTPS using Application Load Balancers (ALBs). The ALBs are configured to use public SSL/TLS certificates.
A Solutions Architect needs to migrate the web applications to a multi-region architecture. All HTTPS services should continue to work without interruption.
Which approach meets these requirements?

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

Answer: D

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

NEW QUESTION 76
A company has more than 100 AWS accounts, with one VPC per account, that need outbound HTTPS connectivity to the internet. The current design contains one NAT gateway per Availability Zone (AZ) in each VPC. To reduce costs and obtain information about outbound traffic, management has asked for a new architecture for internet access.
Which solution will meet the current needs, and continue to grow as new accounts are provisioned, while reducing costs?

A. Create a transit VPC across two AZs using a third-party routing applianc
B. Create a VPN connection to each VP
C. Default route internet traffic to the transit VPC.
D. Create multiple hosted-private AWS Direct Connect VIFs, one per account, each with a Direct Connect gatewa
E. Default route internet traffic back to an on-premises router to route to the internet.
F. Create a central VPC for outbound internet traffi
G. Use VPC peering to default route to a set of redundant NAT gateway in the central VPC.
H. Create a proxy fleet in a central VPC accoun
I. Create an AWS PrivateLink endpoint service in the central VP
J. Use PrivateLink interface for internet connectivity through the proxy fleet.

Answer: D

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

NEW QUESTION 80
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 84
A company has a large on-premises Apache Hadoop cluster with a 20 PB HDFS database. The cluster is growing every quarter by roughly 200 instances and 1 PB. The company’s goals are to enable resiliency for its Hadoop data, limit the impact of losing cluster nodes, and significantly reduce costs. The current cluster runs 24/7 and supports a variety of analysis workloads, including interactive queries and batch processing.
Which solution would meet these requirements with the LEAST expense and down time?

A. Use AWS Snowmobile to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster initially sized to handle the interactive

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

Answer: A

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

NEW QUESTION 89
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 94
An organization has a write-intensive mobile application that uses Amazon API Gateway, AWS Lambda, and Amazon DynamoDB. The application has scaled well, however, costs have increased exponentially because of higher than anticipated Lambda costs. The application's use is unpredictable, but there has been a steady 20% increase in utilization every month.
While monitoring the current Lambda functions, the Solutions Architect notices that the execution-time averages 4.5 minutes. Most of the wait time is the result of a high-latency network call to a 3-TB MySQL database server that is on-premises. A VPN is used to connect to the VPC, so the Lambda functions have been configured with a five-minute timeout.
How can the Solutions Architect reduce the cost of the current architecture?

A. Replace the VPN with AWS Direct Connect to reduce the network latency to the on-premises MySQL database.Enable local caching in the mobile application to reduce the Lambda function invocation calls.Monitor the Lambda function performance; gradually adjust the timeout and memory properties to lower values while maintaining an acceptable execution time.Offload the frequently accessed records from DynamoDB to Amazon ElastiCache.
B. Replace the VPN with AWS Direct Connect to reduce the network latency to the on-premises MySQL database.Cache the API Gateway results to Amazon CloudFront.Use Amazon EC2 Reserved Instances instead of Lambda.Enable Auto Scaling on EC2, and use Spot Instances during peak times.Enable DynamoDB Auto Scaling to manage target utilization.
C. Migrate the MySQL database server into a Multi-AZ Amazon RDS for MySQL.Enable caching of the Amazon API Gateway results in Amazon CloudFront to reduce the number of Lambda function invocations.Monitor the Lambda function performance; gradually adjust the timeout and memory properties to lower values while maintaining an acceptable execution time.Enable DynamoDB Accelerator for frequently accessed records, and enable the DynamoDB Auto Scaling feature.
D. Migrate the MySQL database server into a Multi-AZ Amazon RDS for MySQL.Enable API caching on API Gateway to reduce the number of Lambda function invocations.Continue to monitor the AWS Lambda function performance; gradually adjust the timeout and memory properties to lower values while maintaining an acceptable execution time.Enable Auto Scaling in DynamoDB.

Answer: D

NEW QUESTION 95
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 100
A company needs to run a software package that has a license that must be run on the same physical host for the duration of its use. The software package is only going to be used for 90 days. The company requires patching and restarting of all instances every 30 days.
How can these requirements be met using AWS?

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

Answer: B

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

NEW QUESTION 102
A company has decided to move some workloads onto AWS to create a grid environment to run market analytics. The grid will consist of many similar instances, spun-up by a job-scheduling function. Each time a large analytics workload is completed, a new VPC is deployed along with job scheduler and grid nodes. Multiple grids could be running in parallel.
Key requirements are:

- Grid instances must communicate with Amazon S3 retrieve data to be processed.
- Grid instances must communicate with Amazon DynamoDB to track intermediate data,
- The job scheduler need only to communicate with the Amazon EC2 API to start new grid nodes.

A key requirement is that the environment has no access to the internet, either directly or via the on-premises proxy. However, the application needs to be able to seamlessly communicate to Amazon S3, Amazon DynamoDB, and Amazon EC2 API, without the need for reconfiguration for each new deployment.
Which of the following should the Solutions Architect do to achieve this target architecture? (Choose three.)

A. Enable VPC endpoints for Amazon S3 and DynamoDB.
B. Disable Private DNS Name Support.
C. Configure the application on the grid instances to use the private DNS name of the Amazon S3 endpoint.
D. Populate the on-premises DNS server with the private IP addresses of the EC2 endpoint.
E. Enable an interface VPC endpoint for EC2.
F. Configure Amazon S3 endpoint policy to permit access only from the grid nodes.

Answer: ACE

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/connect-s3-vpc-endpoint/ https://docs.aws.amazon.com/vpc/latest/userguide/vpce-interface.html

NEW QUESTION 103
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 108
A financial services company logs personality identifiable information to its application logs stored in Amazon S3. Due to regulatory compliance requirements, the log files must be encrypted at rest. The Security team has mandated that the company’s on-premises hardware security modules (HSMs) be used to generate the CMK material.
Which steps should the Solution Architected take to meet these requirements?

A. Create an AWS CloudHSM cluste

B. Create a new CMK in AWS KMS using AWS_CloudHSM as the source for the key material and an origin of AWS-CLOUDHS
C. Enable automatic key rotation on the CMK with a duration of 1 yea
D. Configure a bucket policy on the logging bucket the disallow uploads of unencrypted data and requires that the encryption source be AWS KMS.
E. Provision AN AWS Direct Connect connection, ensuring there is no overlap of the RFC 1918 address space between on-premises hardware and the VP
F. Configure an AWS bucket policy on the logging bucket requires all objects to be key material, and create a unique CMK for each logging event.
G. Create a CMK in AWS KMS with no key material and an origin of EXTERNA
H. Import the key material generated from the on-premises HSMs into the CMK using the public key and import token provided by AW
I. Configure a bucket policy on the logging bucket that disallows uploads ofnon-encrypted data and requires that the encryption source be AWS KMS.
J. Create a new CMK in AWS KMS with AWS-provided key material and an origin of AWS-KM
K. Disable this CMK, and overwrite the key material with the material from the on-premises HSM using the public key and import token provided by AWS Re-enable the CM
L. Enable automatic, key rotation on the CMK with a duration of 1 yea
M. Configure a bucket policy on the logging bucket that disallows uploads of non-encrypted data and requires that the encryption source be AWS KMS.

Answer: A

NEW QUESTION 110
A company is running a large application on-premises. Its technology stack consists of Microsoft .NET for the web server platform and Apache Cassandra for the database. The company wants to migrate the application to AWS to improve service reliability. The IT team also wants to reduce the time it spends on capacity management and maintenance of this infrastructure. The Development team is willing and available to make code changes to support the migration.
Which design is the LEAST complex to manage after the migration?

A. Migrate the web servers to Amazon EC2 instances in an Auto Scaling group that is running .NE
B. Migrate the existing Cassandra database to Amazon Aurora with multiple read replicas, and run both in a Multi-AZ mode.
C. Migrate the web servers to an AWS Elastic Beanstalk environment that is running the .NET platform in a Multi-AZ Auto Scaling configuratio
D. Migrate the Cassandra database to Amazon EC2 instances that are running in a Multi-AZ configuration.
E. Migrate the web servers to an AWS Elastic Beanstalk environment that is running the .NET platform in a Multi-AZ Auto Scaling configuratio
F. Migrate the existing Cassandra database to Amazon DynamoDB.
G. Migrate the web servers to Amazon EC2 instances in an Auto Scaling group that is running .NE
H. Migrate the existing Cassandra database to Amazon DynamoDB.

Answer: B

NEW QUESTION 112
A company currently runs a secure application on Amazon EC2 that takes files from on-premises locations through AWS Direct Connect, processes them, and uploads them to a single Amazon S3 bucket. The application uses HTTPS for encryption in transit to Amazon S3, and S3 server-side encryption to encrypt at rest.
Which of the following changes should the Solutions Architect recommend to make this solution more secure without impeding application’s performance?

A. Add a NAT gatewa
B. Update the security groups on the EC2 instance to allow access to and from the S3 IP range onl
C. Configure an S3 bucket policy that allows communication from the NAT gateway’s Elastic IP address only.
D. Add a VPC endpoin
E. Configure endpoint policies on the VPC endpoint to allow access to the required Amazon S3 buckets onl
F. Implement an S3 bucket policy that allows communication from the VPC’s source IP range only.
G. Add a NAT gatewa
H. Update the security groups on the EC2 instance to allow access to and from the S3 IP range onl
I. Configure an S3 bucket policy that allows communication from the source public IP address of the on-premises network only.
J. Add a VPC endpoin
K. Configure endpoint policies on the VPC endpoint to allow access to the required S3 buckets onl
L. Implement an S3 bucket policy that allows communication from the VPC endpoint only.

Answer: D

Explanation:
https://docs.aws.amazon.com/AmazonS3/latest/dev/example-bucket-policies-vpc-endpoint.html

NEW QUESTION 113
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 116
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 120
A company wants to replace its call system with a solution built using AWS managed services. The company call center would like the solution to receive calls, create contact flows, and scale to handle growth projections. The call center would also like the solution to use deep learning capabilities to recognize the intent of the callers and handle basic tasks, reducing the need to speak an agent. The solution should also be able to query business applications and provide relevant information back to calls as requested.
Which services should the Solution Architect use to build this solution? (Choose three.)

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

Answer: BDE

NEW QUESTION 125
A company currently uses a single 1 Gbps AWS Direct Connect connection to establish connectivity between an AWS Region and its data center. The company has five Amazon VPCs, all of which are connected to the data center using the same Direct Connect connection. The Network team is worried about the single point of failure and is interested in improving the redundancy of the connections to AWS while keeping costs to a minimum.
Which solution would improve the redundancy of the connection to AWS while meeting the cost requirements?

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

Answer: B

NEW QUESTION 128
The Security team needs to provide a team of interns with an AWS environment so they can build the serverless video transcoding application. The project will use Amazon S3, AWS Lambda, Amazon API Gateway, Amazon Cognito, Amazon DynamoDB, and Amazon Elastic Transcoder.
The interns should be able to create and configure the necessary resources, but they may not have access to create or modify AWS IAM roles. The Solutions Architect creates a policy and attaches it to the interns’ group.
How should the Security team configure the environment to ensure that the interns are self-sufficient?

A. Create a policy that allows creation of project-related resources onl
B. Create roles with required service permissions, which are assumable by the services.
C. Create a policy that allows creation of all project-related resources, including roles that allow access only to specified resources.
D. Create roles with the required service permissions, which are assumable by the service
E. Have theinterns create and use a bastion host to create the project resources in the project subnet only.
F. Create a policy that allows creation of project-related resources onl
G. Require the interns to raise a request for roles to be created with the Security tea
H. The interns will provide the requirements for the permissions to be set in the role.

Answer: A

NEW QUESTION 130
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 131
A company wants to move a web application to AWS. The application stores session information locally on each web server, which will make auto scaling difficult. As part of the migration, the application will be rewritten to decouple the session data from the web servers. The company requires low latency, scalability, and availability.
Which service will meet the requirements for storing the session information in the MOST cost-effective way?

A. Amazon ElastiCache with the Memcached engine
B. Amazon S3
C. Amazon RDS MySQL
D. Amazon ElastiCache with the Redis engine

Answer: D

Explanation:
https://aws.amazon.com/caching/session-management/ https://aws.amazon.com/elasticache/redis-vs-memcached/

NEW QUESTION 134
A company with multiple accounts is currently using a configuration that does not meet the following security governance policies
• Prevent ingress from port 22 to any Amazon EC2 instance
• Require billing and application tags for resources
• Encrypt all Amazon EBS volumes
A Solutions Architect wants to provide preventive and detective controls including notifications about a specific resource, if there are policy deviations.
Which solution should the Solutions Architect implement?

A. Create an AWS CodeCommit repository containing policy-compliant AWS Cloud Formation templates.Create an AWS Service Catalog portfolio Import the Cloud Formation templates by attaching the CodeCommit repository to the portfolio Restrict users across all accounts to items from the AWSService Catalog portfolio Use AWS Config managed rules to detect deviations from the policie
B. Configure an Amazon CloudWatch Events rule for deviations, and associate a CloudWatch alarm to send notifications when the TriggeredRules metric is greater than zero.
C. Use AWS Service Catalog to build a portfolio with products that are in compliance with the governance policies in a central account Restrict users across all accounts lo AWS Service Catalog products Share a compliant portfolio to other accounts Use AWS Config managed rules to detect deviations from the policies Configure an Amazon CloudWatch Events rule to send a notification when a deviation occurs
D. Implement policy-compliant AWS Cloud Formation templates for each account and ensure that all provisioning is completed by Cloud Formation Configure Amazon Inspector to perform regular checks against resources Perform policy validation and write the assessment output to Amazon CloudWatch Log
E. Create a CloudWatch Logs metric filter to increment a metric when a deviation occurs Configure a CloudWatch alarm to send notifications when the configured metric is greater than zero
F. Restrict users and enforce least privilege access using AWS I A
G. Consolidate all AWS CloudTrail logs into a single account Send the CloudTrail logs to Amazon Elasticsearch Service (Amazon ES). Implement monitoring alerting, and reporting using the Kibana dashboard in Amazon ES and with Amazon SNS.

Answer: C

NEW QUESTION 136
A company has an application written using an in-house software framework. The framework installation takes 30 minutes and is performed with a user data script. Company Developers deploy changes to the application frequently. The framework installation is becoming a bottleneck in this process.
Which of the following would speed up this process?

A. Create a pipeline to build a custom AMI with the framework installed and use this AMI as a baseline for application deployments.
B. Employ a user data script to install the framework but compress the installation files to make them smaller.
C. Create a pipeline to parallelize the installation tasks and call this pipeline from a user data script.
D. Configure an AWS OpsWorks cookbook that installs the framework instead of employing user dat
E. Use this cookbook as a base for all deployments.

Answer: A

Explanation:
https://aws.amazon.com/codepipeline/features/?nc=sn&loc=2

NEW QUESTION 137
A company that is new to AWS reports it has exhausted its service limits across several accounts that are on the Basic Support plan. The company would like to prevent this from happening in the future.
What is the MOST efficient way of monitoring and managing all service limits in the company's accounts?

A. Use Amazon CloudWatch and AWS Lambda to periodically calculate the limits across all linked accounts using AWS Trusted Advisor, provide notifications using Amazon SNS if the limits are close to exceeding the threshold.
B. Reach out to AWS Support to proactively increase the limits across all account
C. That way, the customer avoids creating and managing infrastructure just to raise the service limits.
D. Use Amazon CloudWatch and AWS Lambda to periodically calculate the limits across all linked accounts using AWS Trusted Advisor, programmatically increase the limits that are close to exceeding the threshold.
E. Use Amazon CloudWatch and AWS Lambda to periodically calculate the limits across all linked accounts using AWS Trusted Advisor, and use Amazon SNS for notifications if a limit is close to exceeding the threshol
F. Ensure that the accounts are using the AWS Business Support plan at a minimum.

Answer: D

Explanation:
https://github.com/awslabs/aws-limit-monitor https://aws.amazon.com/solutions/limit-monitor/

NEW QUESTION 138
A company must deploy multiple independent instances of an application. The front-end application is internet accessible. However, corporate policy stipulates that the backends are to be isolated from each other and the internet, yet accessible from a centralized administration server. The application setup should be automated to minimize the opportunity for mistakes as new instances are deployed.
Which option meets the requirements and MINIMIZES costs?

A. Use an AWS CloudFormation template to create identical IAM roles for each regio
B. Use AWS CloudFormation StackSets to deploy each application instance by using parameters to customize for each instance, and use security groups to isolate each instance while permitting access to the central server.
C. Create each instance of the application IAM roles and resources in separate accounts by using AWS CloudFormation StackSet
D. Include a VPN connection to the VPN gateway of the central administration server.
E. Duplicate the application IAM roles and resources in separate accounts by using a single CloudFormation templat
F. Include VPC peering to connect the VPC of each application instance to acentral VPC.
G. Use the parameters of the AWS CloudFormation template to customize the deployment into separate account
H. Include a NAT gateway to allow communication back to the central administration server.

Answer: A

NEW QUESTION 141
A company is running multiple applications on Amazon EC2. Each application is deployed and managed by multiple business units. All applications are deployed on a single AWS account but on different virtual private clouds (VPCs). The company uses a separate VPC in the same account for test and development purposes.
Production applications suffered multiple outages when users accidentally terminated and modified resources that belonged to another business unit. A Solutions Architect has been asked to improve the availability of the company applications while allowing the Developers access to the resources they need.
Which option meets the requirements with the LEAST disruption?

A. Create an AWS account for each business uni
B. Move each business unit's instances to its own account and set up a federation to allow users to access their business unit's account.
C. Set up a federation to allow users to use their corporate credentials, and lock the users down to their own VP
D. Use a network ACL to block each VPC from accessing other VPCs.
E. Implement a tagging policy based on business unit
F. Create an IAM policy so that each user can terminate instances belonging to their own business units only.
G. Set up role-based access for each user and provide limited permissions based on individual roles and the services for which each user is responsible.

Answer: C

Explanation:
Principal – Control what the person making the request (the principal) is allowed to do based on the tags that are attached to that person's IAM user or role. To do this, use the aws:PrincipalTag/key-name condition key to specify what tags must be attached to the IAM user or role before the request is allowed.
https://docs.aws.amazon.com/IAM/latest/UserGuide/access_iam-tags.html

NEW QUESTION 143
A bank is re-architecting its mainframe-based credit card approval processing application to a cloud-native application on the AWS cloud.
The new application will receive up to 1,000 requests per second at peak load. There are multiple steps to each transaction, and each step must receive the result of the previous step. The entire request must return an authorization response within less than 2 seconds with zero data loss. Every request must receive a response. The solution must be Payment Card Industry Data Security Standard (PCI DSS)-compliant.
Which option will meet all of the bank's objectives with the LEAST complexity and LOWEST cost while also meeting compliance requirements?

A. Create an Amazon API Gateway to process inbound requests using a single AWS Lambda task that performs multiple steps and returns a JSON object with the approval statu
B. Open a support case to increase the limit for the number of concurrent Lambdas to allow room for bursts of activity due to the new application.
C. Create an Application Load Balancer with an Amazon ECS cluster on Amazon EC2 Dedicated instances in a target group to process incoming request
D. Use Auto Scaling to scale the cluster out/in based on average CPU utilizatio
E. Deploy a web service that processes all of the approval steps and returns a JSON object with the approval status.
F. Deploy the application on Amazon EC2 on Dedicated Instance
G. Use an Elastic Load Balancer in front of a farm of application servers in an Auto Scaling group to handle incoming request
H. Scale out/in based on a custom Amazon CloudWatch metric for the number of inbound requests per second after measuring the capacity of a single instance.
I. Create an Amazon API Gateway to process inbound requests using a series of AWS Lambda processes, each with an Amazon SQS input queu
J. As each step completes, it writes its result to the next step's queu
K. The final step returns a JSON object with the approval statu

L. Open a support case to increase the limit for the number of concurrent Lambdas to allow room for bursts of activity due to the new application.

Answer: B

NEW QUESTION 147

......