Amazon-Web-Services

Exam Questions SAP-C01

AWS Certified Solutions Architect- Professional

About Exambible

Your Partner of IT Exam

Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

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

≫ An application tier that contains 8 m4.xlarge instances

≫ A Classic Load Balancer

≫ Amazon S3 as a persistent data store
After one of the EC2 instances fails, users report very slow processing of their requests. A Solutions Architect must recommend design changes to maximize system reliability. The solution must minimize costs.
What should the Solution Architect recommend?

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

Answer: B

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

https://aws.amazon.com/blogs/aws/elb-connection-draining-remove-instances-from-service-with-care/

NEW QUESTION 4
A company wants to follow its website on AWS using serverless architecture design patterns for global customers. The company has outlined its requirements as follow:

> The website should be responsive.

> The website should offer minimal latency.

> The website should be highly available.

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

> There should be baseline DDoS protections for spikes in traffic.
How can the design requirements be met?

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

Answer: C

NEW QUESTION 5
A company is designing a new highly available web application on AWS. The application requires consistent and reliable connectivity from the application servers in AWS to a backend REST API hosted in the company's on-premises environment. The backend connection between AWS and on-premises will be routed over an AWS Direct Connect connection through a private virtual interface. Amazon Route 53 will be used to manage private DNS records for the application to resolve the IP address on the backend REST API.
Which design would provide a reliable connection to the backend API?

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

Answer: A

NEW QUESTION 6
An organization has recently grown through acquisitions. Two of the purchased companies use the same IP CIDR range. There is a new short-term requirement to allow AnyCompany A (VPC-A) to communicate with a server that has the IP address 10.0.0.77 in AnyCompany B (VPC-B). AnyCompany A must also communicate with all resources in AnyCompany C (VPC-C). The Network team has created the VPC peer links, but it is having issues with communications between VPC-A and VPC-B. After an investigation, the team believes that the routing tables in the VPCs are incorrect.
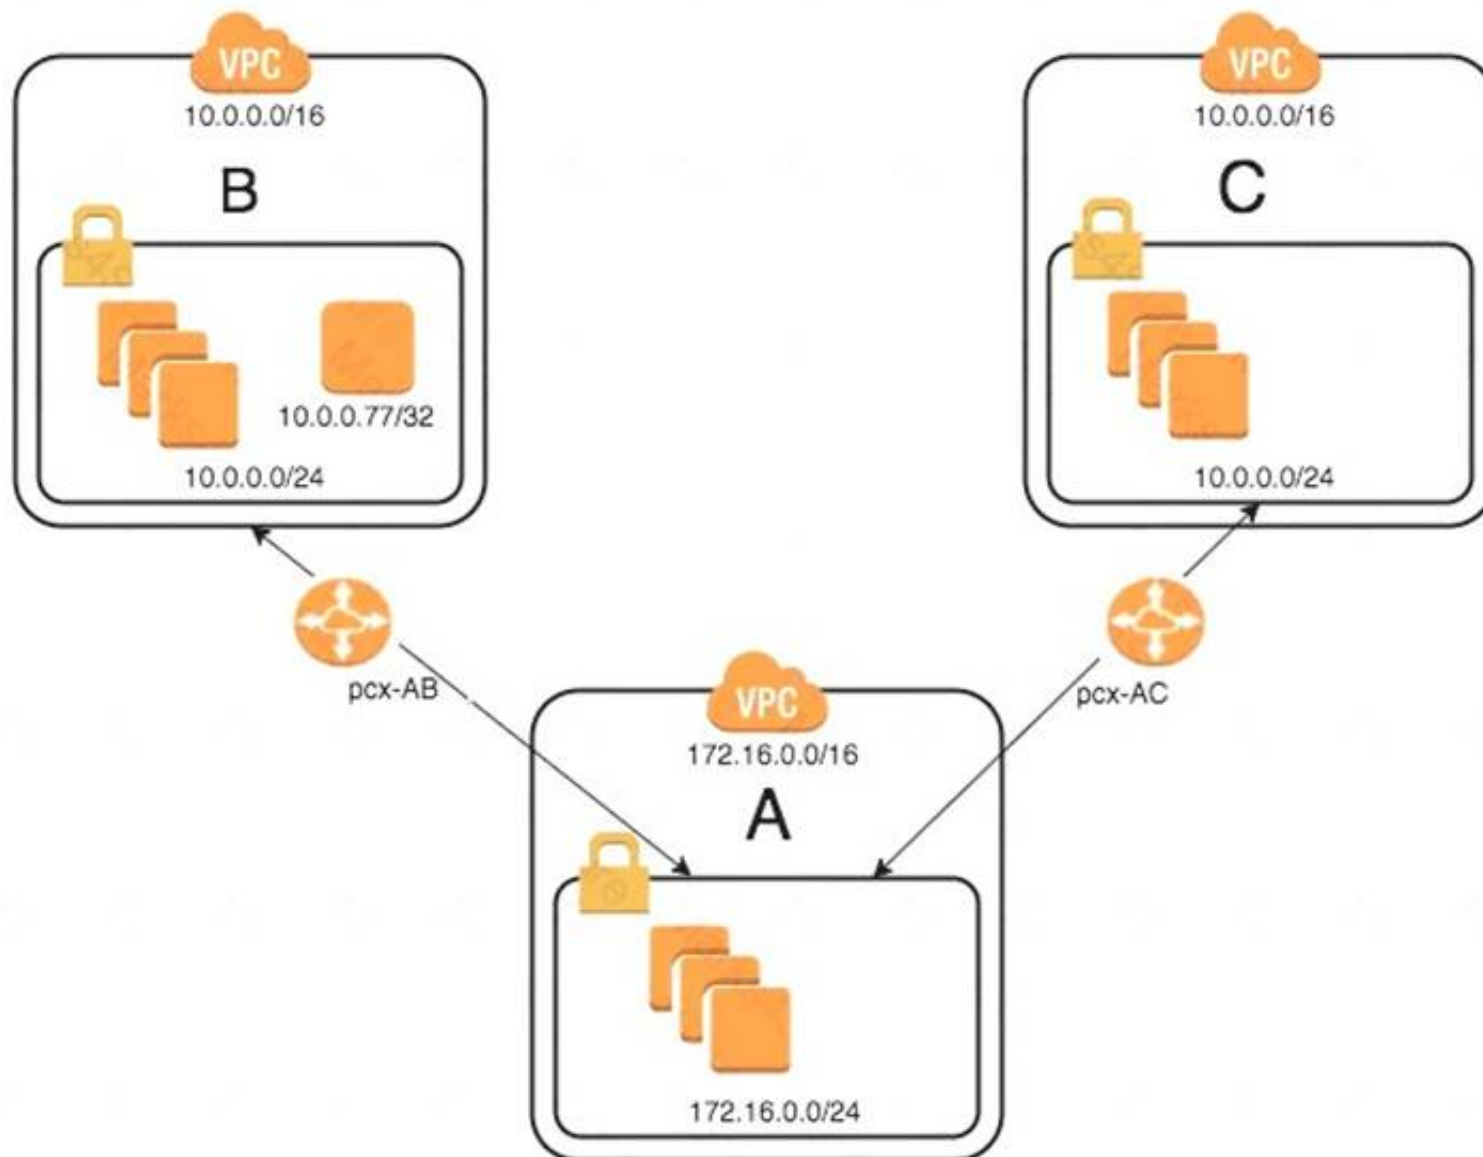
What configuration will allow AnyCompany A to communicate with AnyCompany C in addition to the database in AnyCompany B?

A. On VPC-A, create a static route for the VPC-B CIDR range (10.0.0.0/24) across VPC peerpcx-AB.Create a static route of 10.0.0.0/16 across VPC peer pcx-AC.On VPC-B, create a static route for VPC-A CIDR (172.16.0.0/24) on peer pcx-AB.On VPC-C, create a static route for VPC-A CIDR (172.16.0.0/24) across peer pcx-AC.
B. On VPC-A, enable dynamic route propagation on pcx-AB and pcx-AC.On VPC-B, enable dynamic route propagation and use security groups to allow only the IP address 10.0.0.77/32 on VPC peer pcx-AB.On VPC-C, enable dynamic route propagation with VPC-A on peer pcx-AC.
C. On VPC-A, create network access control lists that block the IP address 10.0.0.77/32 on VPC peerpcx-AC.On VPC-A, create a static route for VPC-B CIDR (10.0.0.0/24) on pcx-AB and a static route for VPC-C CIDR (10.0.0.0/24) on pcx-AC.On VPC-B, create a static route for VPC-A CIDR (172.16.0.0/24) across peer pcx-AB.On VPC-C, create a static route for VPC-A CIDR (172.16.0.0/24) across peer pcx-AC.
D. On VPC-A, create a static route for the VPC-B CIDR (10.0.0.77/32) database across VPC peerpcx-AB.Create a static route for the VPC-C CIDR on VPC peer pcx-AC.On VPC-B, create a static route for VPC-A CIDR (172.16.0.0/24) on peer pcx-AB.On VPC-C, create a static route for VPC-A CIDR (172.16.0.0/24) across peer pcx-AC.

Answer: D

NEW QUESTION 7
A company uses an Amazon EMR cluster to process data once a day. The raw data comes from Amazon S3, and the resulting processed data is also stored in Amazon S3. The processing must complete within 4 hours; currently, it only takes 3 hours. However, the processing time is taking 5 to 10 minutes. longer each week due to an increasing volume of raw data.
The team is also concerned about rising costs as the compute capacity increases. The EMR cluster is currently running on three m3.xlarge instances (one master and two core nodes).
Which of the following solutions will reduce costs related to the increasing compute needs?

A. Add additional task nodes, but have the team purchase an all-upfront convertible Reserved Instance for each additional nod e to offset the costs.
B. Add additional task nodes, but use instance fleets with the master node in on-Demand mode and a mix of On-Demand and Spot Instances for the core and task node
C. Purchase a scheduled Reserved Instances for the master node.
D. Add additional task nodes, but use instance fleets with the master node in Spot mode and a mix of On-Demand and Spot Instances for the core and task node
E. Purchase enough scheduled Reserved Instances to offset the cost of running any On-Demand instances.
F. Add additional task nodes, but use instance fleets with the master node in On-Demand mode and a mix of On-Demand and Spot Instances for the core and task node
G. Purchase a standard all-upfront Reserved Instance for the master node.

Answer: B

NEW QUESTION 8
A company is operating a large customer service call center, and stores and processes call recordings with a custom application Approximately 2% of the call recording are transcribed by an offshore team for quality assurance purposes. These recordings take days. The company uses Linux servers for processing the call recording and managing the transcription queue. There is also a web application for the quality assurance staff to review and score call recordings.
The company plans to migrate the system to AWS to reduce storage costs and the time required to transcribe calls.
Which set of actions should be taken to meet the company's objectives?

A. Upload the call recording to Amazon S3 from the call cente
B. Set up an S3 lifecycle policy to move the call recordings to Amazon S3 Glacier after 90 day
C. Use an AWS Lambda trigger to transcribe the call recordings with Amazon Transcrib

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

Answer: B

NEW QUESTION 9
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 10
A Development team is deploying new APIs as serverless applications within a company. The team is currently using the AWS Management Console to provision Amazon API Gateway, AWS Lambda, and Amazon DynamoDB resources. A Solutions Architect has been tasked with automating the future deployments of these serverless APIs.
How can this be accomplished?

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

Answer: B

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

NEW QUESTION 10
A company is running a .NET three-tier web application on AWS. The team currently uses XL storage optimized instances to store serve the website's image and video files on local instance storage. The company has encountered issues with data loss from replication and instance failures. The Solutions Architect has been asked to redesign this application to improve its reliability while keeping costs low.
Which solution will meet these requirements?

A. Set up a new Amazon EFS share, move all image and video files to this share, and then attach this new drive as a mount point to all existing server
B. Create an Elastic Load Balancer with Auto Scaling general purpose instance
C. Enable Amazon CloudFront to the Elastic Load Balance
D. Enable Cost Explorer and use AWS Trusted advisor checks to continue monitoring the environment for future savings.
E. Implement Auto Scaling with general purpose instance types and an Elastic Load Balance
F. Enable an Amazon CloudFront distribution to Amazon S3 and move images and video files to Amazon S3. Reserve general purpose instances to meet base performance requirement
G. Use Cost Explorer and AWSTrusted Advisor checks to continue monitoring the environment for future savings.
H. Move the entire website to Amazon S3 using the S3 website hosting featur
I. Remove all the web servers and have Amazon S3 communicate directly with the application servers in Amazon VPC.
J. Use AWS Elastic Beanstalk to deploy the .NET applicatio
K. Move all images and video files to Amazon EF
L. Create an Amazon CloudFront distribution that points to the EFS shar
M. Reserve the m4.4xl instances needed to meet base performance requirements.

Answer: B

NEW QUESTION 12
A three-tier web application runs on Amazon EC2 instances. Cron daemons are used to trigger scripts that collect the web server, application, and database logs and send them to a centralized location every hour. Occasionally, scaling events or unplanned outages have caused the instances to stop before the latest logs were collected, and the log files were lost.
Which of the following options is the MOST reliable way of collecting and preserving the log files?

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

Answer: C

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

NEW QUESTION 14
A large multinational company runs a timesheet application on AWS that is used by staff across the world. The application runs on Amazon EC2 instances in an Auto Scaling group behind an Elastic Load Balancing (ELB) load balancer, and stores in an Amazon RDS MySQL Multi-AZ database instance.
The CFO is concerned about the impact on the business if the application is not available. The application must not be down for more than two hours, but the solution must be as cost-effective as possible.
How should the Solutions Architect meet the CFO's requirements while minimizing data loss?

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

Answer: D

NEW QUESTION 19
A company wants to ensure that the workloads for each of its business units have complete autonomy and a minimal blast radius in AWS. The Security team must be able to control access to the resources and services in the account to ensure that particular services are not used by the business units.
How can a Solutions Architect achieve the isolation requirements?

A. Create individual accounts for each business unit and add the account to an OU in AWS Organizations.Modify the OU to ensure that the particular services are blocke
B. Federate each account with an IdP, and create separate roles for the business units and the Security team.
C. Create individual accounts for each business uni
D. Federate each account with an IdP and create separate roles and policies for business units and the Security team.
E. Create one shared account for the entire compan

F. Create separate VPCs for each business uni
G. Create individual IAM policies and resource tags for each business uni
H. Federate each account with an IdP, and create separate roles for the business units and the Security team.
I. Create one shared account for the entire compan
J. Create individual IAM policies and resource tags for each business uni
K. Federate the account with an IdP, and create separate roles for the business units and the Security team.

Answer: A

NEW QUESTION 20
A company is using AWS for production and development workloads. Each business unit has its own AWS account for production, and a separate AWS account to develop and deploy its applications. The Information Security department has introduced new security policies that limit access for terminating certain Amazon ECs instances in all accounts to a small group of individuals from the Security team.
How can the Solutions Architect meet these requirements?

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

Answer: B

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

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

Answer: C

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

NEW QUESTION 27
A company wants to manage the costs associated with a group of 20 applications that are critical, by migrating to AWS. The applications are a mix of Java and Node.js spread across different instance clusters. The company wants to minimize costs while standardizing by using a single deployment methodology. Most of the applications are part of month-end processing routines with a small number of concurrent users, but they are occasionally run at other times. Average application memory consumption is less than 1 GB, though some applications use as much as 2.5 GB of memory during peak processing. The most important application in the group is a billing report written in Java that accesses multiple data sources and often for several hours.
Which is the MOST cost-effective solution?

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

Answer: C

NEW QUESTION 30
An advisory firm is creating a secure data analytics solution for its regulated financial services users Users will upload their raw data to an Amazon 53 bucket, where they have PutObject permissions only Data will be analyzed by applications running on an Amazon EMR cluster launched in a VPC The firm requires that the environment be isolated from the internet All data at rest must be encrypted using keys controlled by the firm

Which combination of actions should the Solutions Architect take to meet the user's security requirements? (Select TWO)

A. Launch the Amazon EMR cluster m a private subnet configured to use an AWS KMS CMK for at-rest encryption Configure a gateway VPC endpoint (or Amazon S3 and an interlace VPC endpoint for AWS KMS
B. Launch the Amazon EMR cluster in a private subnet configured to use an AWS KMS CMK for at-rest encryption Configure a gateway VPC endpomint for Amazon S3 and a NAT gateway to access AWS KMS
C. Launch the Amazon EMR cluster in a private subnet configured to use an AWS CloudHSM appliance for at-rest encryption Configure a gateway VPC endpoint for Amazon S3 and an interface VPC endpoint for CloudHSM
D. Configure the S3 endpoint policies to permit access to the necessary data buckets only
E. Configure the S3 bucket polices lo permit access using an aws sourceVpce condition lo match the S3 endpoint ID

Answer: AC

NEW QUESTION 33
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 36
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 40
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 44
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 46
A company is creating an account strategy so that they can begin using AWS. The Security team will provide each team with the permissions they need to follow the principle or least privileged access. Teams would like to keep their resources isolated from other groups, and the Finance team would like each team's resource usage separated for billing purposes.
Which account creation process meets these requirements and allows for changes?

A. Create a new AWS Organizations accoun
B. Create groups in Active Directory and assign them to roles in AWS to grant federated acces
C. Require each team to tag their resources, and separate bills based on tag
D. Control access to resources through IAM granting the minimally required privilege.
E. Create individual accounts for each tea
F. Assign the security as the master account, and enable consolidated billing for all other account
G. Create a cross-account role for security to manage accounts, and send logs to a bucket in the security account.
H. Create a new AWS account, and use AWS Service Catalog to provide teams with the required resources.Implement a third-party billing to provide the Finance team with the resource use for each team based on taggin
I. Isolate resources using IAM to avoid account spraw
J. Security will control and monitor logs and permissions.
K. Create a master account for billing using Organizations, and create each team's account from that master accoun
L. Create a security account for logs and cross-account acces
M. Apply service control policies on each account, and grant the Security team cross-account access to all account
N. Security will create IAM policies for each account to maintain least privilege access.

Answer: B

NEW QUESTION 50
A company runs a Windows Server host in a public subnet that is configured to allow a team of administrators to connect over RDP to troubleshoot issues with hosts in a private subnet. The host must be available at all times outside of a scheduled maintenance window, and needs to receive the latest operating system updates within 3 days of release.
What should be done to manage the host with the LEAST amount of administrative effort?

A. Run the host in a single-instance AWS Elastic Beanstalk environmen
B. Configure the environment with a custom AMI to use a hardened machine image from AWS Marketplac
C. Apply system updates with AWS Systems Manager Patch Manager.
D. Run the host on AWS WorkSpace
E. Use Amazon WorkSpaces Application Manager (WAM) to harden the hos
F. Configure Windows automatic updates to occur every 3 days.
G. Run the host in an Auto Scaling group with a minimum and maximum instance count of 1. Use a hardened machine image from AWS Marketplac
H. Apply system updates with AWS Systems Manager Patch Manager.
I. Run the host in AWS OpsWorks Stack
J. Use a Chief recipe to harden the AMI during instance launch.Use an AWS Lambda scheduled event to run the Upgrade Operating System stack command to apply system updates.

Answer: B

NEW QUESTION 51
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 54
A company has been using a third-party provider for its content delivery network and recently decided to switch to Amazon CloudFront the Development team wants to maximize performance for the global user base. The company uses a content management system (CMS) that serves both static and dynamic content. The CMS is both md an Application Load Balancer (ALB) which is set as the default origin for the distribution. Static assets are served from an Amazon S3 bucket. The Origin Access Identity (OAI) was created property d the S3 bucket policy has been updated to allow the GetObject action from the OAI, but static assets are receiving a 404 error
Which combination of steps should the Solutions Architect take to fix the error? (Select TWO.)

A. Add another origin to the CloudFront distribution for the static assets
B. Add a path based rule to the ALB to forward requests for the static assets
C. Add an RTMP distribution to allow caching of both static and dynamic content
D. Add a behavior to the CloudFront distribution for the path pattern and the origin of the static assets
E. Add a host header condition to the ALB listener and forward the header from CloudFront to add traffic to the allow list

Answer: AD

NEW QUESTION 57
A Solutions Architect is migrating a 10 TB PostgreSQL database to Amazon RDS for PostgreSQL. The company's internet link is 50 MB with a VPN in the Amazon VPC, and the Solutions Architect needs to migrate the data and synchronize the changes before the cutover. The cutover must take place within an 8-day period.
What is the LEAST complex method of migrating the database securely and reliably?

A. Order an AWS Snowball device and copy the database using the AWS DM
B. When the database is available in Amazon 3, use AWS DMS to load it to Amazon RDS, and configure a job to synchronize changes before the cutover.
C. Create an AWS DMS job to continuously replicate the data from on premises to AW
D. Cutover to Amazon RDS after the data is synchronized.
E. Order an AWS Snowball device and copy a database dump to the devic
F. After the data has been copied to Amazon S3, import it to the Amazon RDS instanc
G. Set up log shipping over a VPN to synchronize changes before the cutover.
H. Order an AWS Snowball device and copy the database by using the AWS Schema Conversion Tool.When the data is available in Amazon S3, use AWS DMS to load it to Amazon RDS, and configure a job to synchronize changes before the cutover.

Answer: B

NEW QUESTION 61

A company is adding a new approved external vendor that only supports IPv6 connectivity. The company's backend systems sit in the private subnet of an Amazon VPC. The company uses a NAT gateway to allow these systems to communicate with external vendors over IPv4. Company policy requires systems that communicate with external vendors use a security group that limits access to only approved external vendors. The virtual private cloud (VPC) uses the default network ACL.

The Systems Operator successfully assigns IPv6 addresses to each of the backend systems. The Systems Operator also updates the outbound security group to include the IPv6 CIDR of the external vendor (destination). The systems within the VPC are able to ping one another successfully over IPv6. However, these systems are unable to communicate with the external vendor.

What changes are required to enable communication with the external vendor?

A. Create an IPv6 NAT instanc
B. Add a route for destination 0.0.0.0/0 pointing to the NAT instance.
C. Enable IPv6 on the NAT gatewa
D. Add a route for destination ::/0 pointing to the NAT gateway.
E. Enable IPv6 on the internet gatewa
F. Add a route for destination 0.0.0.0/0 pointing to the IGW.
G. Create an egress-only internet gatewa
H. Add a route for destination ::/0 pointing to the gateway.

Answer: D

Explanation:
https://docs.aws.amazon.com/vpc/latest/userguide/egress-only-internet-gateway.html

NEW QUESTION 63

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

➤ Aggregate logs using AWS.

➤ Automate log analysis for errors.

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

A. Install Amazon Kinesis Agent on servers, send logs to Amazon Kinesis Data Streams and use Amazon Kinesis Data Analytics to identify errors, create an Amazon CloudWatch alarm to notify the Operations team of errors
B. Install an AWS X-Ray agent on servers, send logs to AWS Lambda and analyze them to identify errors, use Amazon CloudWatch Events to notify the Operations team of errors.
C. Install Logstash on servers, send logs to Amazon S3 and use Amazon Athena to identify errors, use sendmail to notify the Operations team of errors.
D. Install the Amazon CloudWatch agent on servers, send logs to Amazon CloudWatch Logs and use metric filters to identify errors, create a CloudWatch alarm to notify the Operations team of errors.

Answer: A

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

NEW QUESTION 66

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

Which solution meets the company's requirements?

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

Answer: B

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

NEW QUESTION 67

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 71
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 72
A company is currently using AWS CodeCommit for its source control and AWS CodePipeline for continuous integration. The pipeline has a build stage for building the artifacts which is then staged in an Amazon S3 bucket.
The company has identified various improvement opportunities in the existing process, and a Solutions Architect has been given the following requirement:

≫ Create a new pipeline to support feature development

≫ Support feature development without impacting production applications

≫ Incorporate continuous testing with unit tests

≫ Isolate development and production artifacts

≫ Support the capability to merge tested code into production code. How should the Solutions Architect achieve these requirements?

A. Trigger a separate pipeline from CodeCommit feature branche
B. Use AWS CodeBuild for running unit test
C. Use CodeBuild to stage the artifacts within an S3 bucket in a separate testing account.
D. Trigger a separate pipeline from CodeCommit feature branche
E. Use AWS Lambda for running unit test
F. Use AWS CodeDeploy to stage the artifacts within an S3 bucket in a separate testing account.
G. Trigger a separate pipeline from CodeCommit tags Use Jenkins for running unit test
H. Create a stage in the pipeline with S3 as the target for staging the artifacts with an S3 bucket in a separate testing account.
I. Create a separate CodeCommit repository for feature development and use it to trigger the pipelin
J. Use AWS Lambda for running unit test
K. Use AWS CodeBuild to stage the artifacts within different S3 buckets in the same production account.

Answer: A

Explanation:
https://docs.aws.amazon.com/codebuild/latest/userguide/how-to-create-pipeline.html

NEW QUESTION 75
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 76
A company has an application behind a load balancer with enough Amazon EC2 instances to satisfy peak demand. Scripts and third-party deployment solutions are used to configure EC2 instances when demand increases or an instance fails. The team must periodically evaluate the utilization of the instance types to ensure that the correct sizes are deployed.
How can this workload be optimized to meet these requirements?

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

Answer: D

NEW QUESTION 80
A Solutions Architect needs to design a highly available application that will allow authenticated users to stay connected to the application even when there are underlying failures
Which solution will meet these requirements?

A. Deploy the application on Amazon EC2 instances Use Amazon Route 53 to forward requests to the EC2 Instances Use Amazon DynamoDB to save the authenticated connection details
B. Deploy the application on Amazon EC2 instances in an Auto Scaling group Use an internet-facing Application Load Balancer to handle requests Use Amazon DynamoDB to save the authenticated connection details
C. Deploy the application on Amazon EC2 instances in an Auto Scaling group Use an internet-facing Application Load Balancer on the front end Use EC2 instances to save the authenticated connectiondetails
D. Deploy the application on Amazon EC2 instances in an Auto Scaling group Use an internet-facing Application Load Balancer on the front end Use EC2 instances hosting a MySQL database to save the authenticated connection details

Answer: B

NEW QUESTION 84
A media storage application uploads user photos to Amazon S3 for processing. End users are reporting that some uploaded photos are not being processed properly. The Application Developers trace the logs and find that AWS Lambda is experiencing execution issues when thousands of users are on the system simultaneously. Issues are caused by:

> Limits around concurrent executions.

> The performance of Amazon DynamoDB when saving data.
Which actions can be taken to increase the performance and reliability of the application? (Choose two.)

A. Evaluate and adjust the read capacity units (RCUs) for the DynamoDB tables.
B. Evaluate and adjust the write capacity units (WCUs) for the DynamoDB tables.
C. Add an Amazon ElastiCache layer to increase the performance of Lambda functions.
D. Configure a dead letter queue that will reprocess failed or timed-out Lambda functions.
E. Use S3 Transfer Acceleration to provide lower-latency access to end users.

Answer: BD

Explanation:
B:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.h
D: https://aws.amazon.com/blogs/compute/robust-serverless-application-design-with-aws-lambda-dlq/c

NEW QUESTION 86
A company is building an AWS landing zone and has asked a Solutions Architect to design a multi-account access strategy that will allow hundreds of users to use corporate credentials to access the AWS Console. The company is running a Microsoft Active Directory and users will use an AWS Direct Connect connection to connect to AWS. The company also wants to be able to federate to third-party services and providers, including custom applications.
Which solution meets the requirements by using the LEAST amount of management overhead?

A. Connect the Active Directory to AWS by using single sign-on and an Active Directory Federation Services (AD FS) with SAML 2.0, and then configure the identity Provider (IdP) system to use form-based authenticatio
B. Build the AD FS portal page with corporate branding, and integrate third-party applications that support SAML 2.0 as required.
C. Create a two-way Forest trust relationship between the on-premises Active Directory and the AWS Directory Servic
D. Set up AWS Single Sign-On with AWS Organization
E. Use single sign-on integrations for connections with third-party applications.
F. Configure single sign-on by connecting the on-premises Active Directory using the AWS Directory Service AD Connecto
G. Enable federation to the AWS services and accounts by using the IAM applications and services linking functio
H. Leverage third-party single sign-on as needed.
I. Connect the company's Active Directory to AWS by using AD FS and SAML 2.0. Configure the AD FS claim rule to leverage Regex and a common Active Directory naming convention for the security group to allow federation of all AWS account
J. Leverage third-party single sign-on as needed, and add it to the AD FS server.

Answer: D

Explanation:
https://aws.amazon.com/blogs/security/aws-federated-authentication-with-active-directory-federation-services-a

NEW QUESTION 90
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 92
A large global company wants to migrate a stateless mission-critical application to AWS. The application is based on IBM WebSphere (application and integration middleware), IBM MQ (messaging middleware), and IBM DB2 (database software) on a z/OS operating system.
How should the Solutions Architect migrate the application to AWS?

A. Re-host WebSphere-based applications on Amazon EC2 behind a load balancer with Auto Scaling.Re-platform the IBM MQ to an Amazon EC2-based M
B. Re-platform the z/OS-based DB2 to Amazon RDS DB2.
C. Re-host WebSphere-based applications on Amazon EC2 behind a load balancer with Auto Scaling.Re-platform the IBM MQ to an Amazon M
D. Re-platform z/OS-based DB2 to Amazon EC2-based DB2.
E. Orchestrate and deploy the application by using AWS Elastic Beanstal
F. Re-platform the IBM MQ to Amazon SQ
G. Re-platform z/OS-based DB2 to Amazon RDS DB2.
H. Use the AWS Server Migration Service to migrate the IBM WebSphere and IBM DB2 to an Amazon EC2-based solutio
I. Re-platform the IBM MQ to an Amazon MQ.

Answer: B

Explanation:
https://aws.amazon.com/blogs/database/aws-database-migration-service-and-aws-schema-conversion-tool-now-
https://aws.amazon.com/quickstart/architecture/ibm-mq/

NEW QUESTION 94
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 97
......

Relate Links

100% Pass Your SAP-C01 Exam with Exambible Prep Materials

https://www.exambible.com/SAP-C01-exam/

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - https://www.exambible.com/