

Amazon-Web-Services

Exam Questions SAP-C01

AWS Certified Solutions Architect- Professional



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 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 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 in Fargate using Service Auto Scaling 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 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.
- 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 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 redundancy.
- 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 Region.
- E. Modify the application to replicate data to both regions.
- F. Establish a second DX connection for redundancy.
- G. Create an identical DynamoDB table in a second Region.
- H. Enable DynamoDB auto scaling to manage throughput capacity.
- 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 table.
- L. Use AWS Lambda to replicate changes to the second Region.

Answer: A

NEW QUESTION 3

An organization has two Amazon EC2 instances:

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

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 4

A large company has many business units. Each business unit has multiple AWS accounts for different purposes. The CIO of the company sees that each business unit has data that would be useful to share with other parts of the company. In total, there are about 10 PB of data that needs to be shared with users in

1,000 AWS accounts. The data is proprietary, so some of it should only be available to users with specific job types. Some of the data is used for throughput of intensive workloads, such as simulations. The number of accounts changes frequently because of new initiatives, acquisitions, and divestitures. A Solutions Architect has been asked to design a system that will allow for sharing data for use in AWS with all of the employees in the company. Which approach will allow for secure data sharing in a scalable way?

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

Answer: B

NEW QUESTION 5

An on-premises application will be migrated to the cloud. The application consists of a single Elasticsearch virtual machine with data source feeds from local systems that will not be migrated, and a Java web application on Apache Tomcat running on three virtual machines. The Elasticsearch server currently uses 1 TB of storage out of 16 TB available storage, and the web application is updated every 4 months. Multiple users access the web application from the Internet. There is a 10Gbit AWS Direct Connect connection established, and the application can be migrated over a scheduled 48-hour change window. Which strategy will have the LEAST impact on the Operations staff after the migration?

- A. Create an Elasticsearch server on Amazon EC2 right-sized with 2 TB of Amazon EBS and a public AWS Elastic Beanstalk environment for the web application
- B. Pause the data sources, export the Elasticsearch index from on premises, and import into the EC2 Elasticsearch server
- C. Move data source feeds to the new Elasticsearch server and move users to the web application.
- D. Create an Amazon ES cluster for Elasticsearch and a public AWS Elastic Beanstalk environment for the web application
- E. Use AWS DMS to replicate Elasticsearch data
- F. When replication has finished, move data source feeds to the new Amazon ES cluster endpoint and move users to the new web application.
- G. Use the AWS SMS to replicate the virtual machines into AWS
- H. When the migration is complete, pause the data source feeds and start the migrated Elasticsearch and web application instance
- I. Place the web application instances behind a public Elastic Load Balance
- J. Move the data source feeds to the new Elasticsearch server and move users to the new web Application Load Balancer.
- K. Create an Amazon ES cluster for Elasticsearch and a public AWS Elastic Beanstalk environment for the web application
- L. Pause the data source feeds, export the Elasticsearch index from on premises, and import into the Amazon ES cluster
- M. Move the data source feeds to the new Amazon ES cluster endpoint and move users to the new web application.

Answer: D

NEW QUESTION 6

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 node 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 nodes
- C. Purchase a scheduled Reserved Instance 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 nodes
- 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 nodes
- G. Purchase a standard all-upfront Reserved Instance for the master node.

Answer: B

NEW QUESTION 7

A company has a requirement that only allows specially hardened AMIs to be launched into public subnets in a VPC, and for the AMIs to be associated with a specific security group. Allowing non-compliant instances to launch into the public subnet could present a significant security risk if they are allowed to operate. A mapping of approved AMIs to subnets to security groups exists in an Amazon DynamoDB table in the same AWS account. The company created an AWS Lambda function that, when invoked, will terminate a given Amazon EC2 instance if the combination of AMI, subnet, and security group are not approved in the DynamoDB table. What should the Solutions Architect do to MOST quickly mitigate the risk of compliance deviations?

- A. Create an Amazon CloudWatch Events rule that matches each time an EC2 instance is launched using one of the allowed AMIs, and associate it with the Lambda function as the target.

- B. For the Amazon S3 bucket receiving the Aws CloudTrail logs, create an S3 event notification configuration with a filter to match when logs contain the ec2:RunInstances action, and associate it with the Lambda function as the target.
- C. Enable AWS CloudTrail and configure it to stream to an Amazon CloudWatch Logs group.
- D. Create a metric filter in CloudWatch to match when the ec2:RunInstances action occurs, and trigger the Lambda function when the metric is greater than 0.
- E. Create an Amazon CloudWatch Events rule that matches each time an EC2 instance is launched, and associate it with the Lambda function as the target.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-lifecycle.html>

NEW QUESTION 8

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

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

Answer: B

Explanation:

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

NEW QUESTION 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 endpoints.
- B. Add AWS WAF rules to enforce the company's requirements.
- C. Deploy Amazon CloudFront in front of all the endpoints.
- D. The CloudFront distribution provides perimeter protection.
- E. Add AWS Lambda-based automation to provide additional security.
- F. Deploy Amazon CloudFront in front of all the endpoints.
- G. Deploy AWS WAF and AWS Shield Advanced.
- H. Add AWS WAF rules to enforce the company's requirements.
- 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 requirements.
- 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 instance
- B. Configure an IAM role to allow the start and stop of the Elastic Beanstalk environment and attach it to the instance
- C. Create scripts on the instance to start and stop the Elastic Beanstalk environment
- D. Configure cron jobs on the instance to execute the scripts.
- E. Develop AWS Lambda functions to start and stop the Elastic Beanstalk environment
- F. Configure a Lambda execution role granting Elastic Beanstalk environment start/stop permissions, and assign the role to the Lambda function
- G. Configure cron expression Amazon CloudWatch Events rules to trigger the Lambda functions.
- H. Develop an AWS Step Functions state machine with "wait" as its type to control the start and stop time. Use the activity task to start and stop the Elastic Beanstalk environment
- I. Create a role for Step Functions to allow it to start and stop the Elastic Beanstalk environment
- J. Invoke Step Functions daily.
- K. Configure a time-based Auto Scaling group
- L. In the morning, have the Auto Scaling group scale up an Amazon EC2 instance and put the Elastic Beanstalk environment start command in the EC2 instance user data
- M. At the end of the day, scale down the instance number to 0 to terminate the EC2 instance.

Answer: B

Explanation:

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

NEW QUESTION 10

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

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

Answer: C

NEW QUESTION 13

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

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

Answer: A

NEW QUESTION 14

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 one other stack. 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 service.
- B. After the existing templates, use cross-stack references to eliminate passing many parameters to each template.
- C. Call each required stack for the application as a nested stack from the new stack.
- D. Call the newly created service stack from the AWS 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 service.
- F. Create a product for each existing AWS CloudFormation template required to build the service.
- G. Add the products to the portfolio that represents that service in AWS Service Catalog.
- 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 service.
- J. For each existing template, choose AWS CloudFormation as a deployment action.
- K. Add the AWS CloudFormation template to the deployment action.
- L. Ensure that the deployment actions are processed to make sure that dependencies are obeyed.
- M. Use configuration files and scripts to share parameters between the stacks.
- 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 service.
- P. Create a new AWS CloudFormation template for each service.
- Q. After the existing templates, use cross-stack references to eliminate passing many parameters to each template.

- R. Call each required stack for the application as a nested stack from the new service template
- S. Configure AWS Step Functions to call the service template directly
- T. In the AWS Step Functions console, execute the step.
- . Create a new portfolio for the Services in AWS Service Catalog
- . Create a new AWS CloudFormation template for each service
- . After the existing templates to use cross-stack references to eliminate passing many parameters to each template
- . Call each required stack for the application as a nested stack from the new stack
- . Create a product for each application
- . Add the service template to the product
- . Add each new product to the portfolio
- . Deploy the product from the portfolio to deploy the service with the necessary parameters only to start the deployment.

Answer: AE

NEW QUESTION 19

A company is running an email application across multiple AWS Regions. The company uses Ohio (us-east-2) as the primary Region and Northern Virginia (us-east-1) as the Disaster Recovery (DR) Region. The data is continuously replicated from the primary Region to the DR Region by a single instance on the public subnet in both Regions. The replication messages between the Regions have a significant backlog during certain times of the day. The backlog clears on its own after a short time, but it affects the application's RPO.

Which of the following solutions should help remediate this performance problem? (Select TWO)

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

Answer: AC

NEW QUESTION 22

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 account
- 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 27

A company is migrating an application to AWS. It wants to use fully managed services as much as possible during the migration. The company needs to store large, important documents within the application with the following requirements:

- The data must be highly durable and available.
- The data must always be encrypted at rest and in transit.
- The encryption key must be managed by the company and rotated periodically. Which of the following solutions should the Solutions Architect recommend?

- A. Deploy the storage gateway to AWS in file gateway mode
- B. Use Amazon EBS volume encryption using an AWS KMS key to encrypt the storage gateway volumes.
- C. Use Amazon S3 with a bucket policy to enforce HTTPS for connections to the bucket and to enforce server-side encryption and AWS KMS for object encryption.
- D. Use Amazon DynamoDB with SSL to connect to DynamoDB
- E. Use an AWS KMS key to encrypt DynamoDB objects at rest.
- F. Deploy instances with Amazon EBS volumes attached to store this data
- G. Use EBS volume encryption using an AWS KMS key to encrypt the data.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/security/how-to-use-bucket-policies-and-apply-defense-in-depth-to-help-secure-y>

NEW QUESTION 28

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

- A. Create a new IAM policy that allows access to those EC2 instances only for the Security team
- B. Apply this policy to the AWS Organizations master account.
- C. Create a new tag-based IAM policy that allows access to these EC2 instances only for the Security team. Tag the instances appropriately, and apply this policy in each account.
- D. Create an organizational unit under AWS Organization
- E. Move all the accounts into this organizational unit and use SCP to apply a whitelist policy to allow access to these EC2 instances for the Security team only.
- F. Set up SAML federation for all accounts in 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 32

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 routing
- 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 region
- 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 region
- 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 layer
- I. Use Amazon DynamoDB global tables as the data layer with Amazon DynamoDB Accelerator (DAX) for caching.
- J. Use Amazon Route 53 geolocation-based routing
- 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 region
- 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 34

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

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

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

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

Answer: C

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-https-connection-fails/>

NEW QUESTION 40

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 41

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

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

Answer: BD

Explanation:

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

NEW QUESTION 46

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 49

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 of 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 account
- B. Create groups in Active Directory and assign them to roles in AWS to grant federated access
- 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 team
- F. Assign the security as the master account, and enable consolidated billing for all other accounts
- 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 tagging
- I. Isolate resources using IAM to avoid account sprawl
- 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 account
- L. Create a security account for logs and cross-account access
- M. Apply service control policies on each account, and grant the Security team cross-account access to all accounts
- N. Security will create IAM policies for each account to maintain least privilege access.

Answer: B

NEW QUESTION 51

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

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

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

Answer: B

NEW QUESTION 53

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

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

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 56

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 59

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 61

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

What is the MOST cost-effective solution?

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

Answer: C

NEW QUESTION 71

A company has more than 100 AWS accounts, with one VPC per account, that need outbound HTTPS connectivity to the internet. The current design contains one

NAT gateway per Availability Zone (AZ) in each VPC. To reduce costs and obtain information about outbound traffic, management has asked for a new architecture for internet access.

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

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

Answer: D

Explanation:

user proxy fleet over PrivateLink. As explained in this AWS website:

<https://aws.amazon.com/blogs/networking-and-content-delivery/how-to-use-aws-privatelink-to-secure-and-scale>

NEW QUESTION 73

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

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

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

Answer: A

Explanation:

Q: How should I choose between Snowmobile and Snowball?

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

NEW QUESTION 76

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

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

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

Answer: BDE

NEW QUESTION 78

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

Which solution will meet the CISO's requirements?

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

Answer: A

Explanation:

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

NEW QUESTION 80

A Solution Architect is designing a deployment strategy for an application tier and has the following requirements.

- * The application code will need a 500 MB 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 unpatches
- 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 patches
- E. Update the Auto Scaling group to use the patches AMI and replace existing unpatches and the application code as a batch job every night
- 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 count
- H. Configure an Amazon EC2 data script to download the data from Amazon S3 install OS patches with AWS system Manager when they are released
- 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 each Amazon-provided AMI release
- K. Use AWS CodeDeploy to push the application code to the instances.

Answer: C

NEW QUESTION 81

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

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

Answer: B

NEW QUESTION 85

A retail company is running an application that stores invoice files in Amazon S3 bucket and metadata about the files in an Amazon DynamoDB table. 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 table
- C. Set up S3 cross-region replication from us-east-1 to eu-west-1. Set up MFA delete on 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 bucket
- 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 to eu-west-1.

Answer: A

Explanation:

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

NEW QUESTION 88

An online retailer needs to regularly process large product catalogs, which are handled in batches. These are sent out to be processed by people using the Amazon Mechanical Turk service, but the retailer has asked its Solutions Architect to design a workflow orchestration system that allows it to handle multiple concurrent Mechanical Turk operations, deal with the result assessment process, and reprocess failures. Which of the following options gives the retailer the ability to interrogate the state of every workflow with the LEAST amount of implementation effort?

- A. Trigger Amazon CloudWatch alarms based upon message visibility in multiple Amazon SQS queues (one queue per workflow stage) and send messages via Amazon SNS to trigger AWS Lambda functions to process the next step
- B. Use Amazon ES and Kibana to visualize Lambda processing logs to see the workflow states.

- C. Hold workflow information in an Amazon RDS instance with AWS Lambda functions polling RDS for status change
- D. Worker Lambda functions then process the next workflow step
- E. Amazon QuickSight will visualize workflow states directly out of Amazon RDS.
- F. Build the workflow in AWS Step Functions, using it to orchestrate multiple concurrent workflow
- G. The status of each workflow can be visualized in the AWS Management Console, and historical data can be written to Amazon S3 and visualized using Amazon QuickSight.
- H. Use Amazon SWF to create a workflow that handles a single batch of catalog records with multiple worker tasks to extract the data, transform it, and send it through Mechanical Tur
- I. Use Amazon ES and Kibana to visualize AWS Lambda processing logs to see the workflow states.

Answer: C

Explanation:

AWS Step Functions is a fully managed service that makes it easy to coordinate the components of distributed applications and microservices using visual workflows. Instead of writing a Decider program, you define state machines in JSON. AWS customers should consider using Step Functions for new applications. If Step Functions does not fit your needs, then you should consider Amazon Simple Workflow (SWF). Amazon SWF provides you complete control over your orchestration logic, but increases the complexity of developing applications. You may write decider programs in the programming language of your choice, or you may use the Flow framework to use programming constructs that structure asynchronous interactions for you. AWS will continue to provide the Amazon SWF service, Flow framework, and support all Amazon SWF customers. <https://aws.amazon.com/swf/faqs/>

NEW QUESTION 89

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 Architect take to meet these requirements?

- A. Create an AWS CloudHSM cluster
- B. Create a new CMK in AWS KMS using AWS_CloudHSM as the source for the key material and an origin of AWS-CLOUDHSM
- C. Enable automatic key rotation on the CMK with a duration of 1 year
- D. Configure a bucket policy on the logging bucket that disallows 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 VPC
- F. Configure an AWS bucket policy on the logging bucket that 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 EXTERNAL
- H. Import the key material generated from the on-premises HSMs into the CMK using the public key and import token provided by AWS
- I. Configure a bucket policy on the logging bucket that disallows uploads of non-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 CMK
- L. Enable automatic key rotation on the CMK with a duration of 1 year
- 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 92

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

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

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

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

Answer: ABD

NEW QUESTION 93

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

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

- A. Amazon Rekognition to identify who is calling.
- B. Amazon Connect to create a cloud-based contact center.
- C. Amazon Alexa for Business to build conversational interface.
- D. AWS Lambda to integrate with internal systems.

- E. Amazon Lex to recognize the intent of the caller.
- F. Amazon SQS to add incoming callers to a queue.

Answer: BDE

NEW QUESTION 97

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 connection details
- 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 101

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 account
- B. Create users and groups in the identity account
- 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 team
- F. Set a strong IAM password policy on each account
- 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 account
- 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 103

A company has an application that generates a weather forecast that is updated every 15 minutes with an output resolution of 1 billion unique positions, each approximately 20 bytes in size (20 Gigabytes per forecast). Every hour, the forecast data is globally accessed approximately 5 million times (1,400 requests per second), and up to 10 times more during weather events. The forecast data is overwritten every update. Users of the current weather forecast application expect responses to queries to be returned in less than two seconds for each request.

Which design meets the required request rate and response time?

- A. Store forecast locations in an Amazon ES cluster
- B. Use an Amazon CloudFront distribution targeting an Amazon API Gateway endpoint with AWS Lambda functions responding to queries as the origin
- C. Enable API caching on the API Gateway stage with a cache-control timeout set for 15 minutes.
- D. Store forecast locations in an Amazon EFS volume
- E. Create an Amazon CloudFront distribution that targets an Elastic Load Balancing group of an Auto Scaling fleet of Amazon EC2 instances that have mounted the Amazon EFS volume
- F. Set the cache-control timeout for 15 minutes in the CloudFront distribution.
- G. Store forecast locations in an Amazon ES cluster
- H. Use an Amazon CloudFront distribution targeting an API Gateway endpoint with AWS Lambda functions responding to queries as the origin
- I. Create an Amazon Lambda@Edge function that caches the data locally at edge locations for 15 minutes.
- J. Store forecast locations in an Amazon S3 as individual objects
- K. Create an Amazon CloudFront distribution targeting an Elastic Load Balancing group of an Auto Scaling fleet of EC2 instances, querying the origin of the S3 objects
- L. Set the cache-control timeout for 15 minutes in the CloudFront distribution.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/networking-and-content-delivery/lambdaedge-design-best-practices/>

NEW QUESTION 105

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 AWS Service Catalog portfolio. Use AWS Config managed rules to detect deviations from the policies.
- 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 to 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 CloudFormation. 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 IAM.
- 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 109

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 unit.
- 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 VPC.
- 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 114

A company collects a steady stream of 10 million data records from 100,000 sources each day. These records are written to an Amazon RDS MySQL DB. A query must produce the daily average of a data source over the past 30 days. There are twice as many reads as writes. Queries to the collected data are for one source ID at a time.

How can the Solutions Architect improve the reliability and cost effectiveness of this solution?

- A. Use Amazon Aurora with MySQL in a Multi-AZ mode.
- B. Use four additional read replicas.
- C. Use Amazon DynamoDB with the source ID as the partition key and the timestamp as the sort key.
- D. Use a Time to Live (TTL) to delete data after 30 days.
- E. Use Amazon DynamoDB with the source ID as the partition key.
- F. Use a different table each day.
- G. Ingest data into Amazon Kinesis using a retention period of 30 days.
- H. Use AWS Lambda to write data records to Amazon ElastiCache for read access.

Answer: B

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html>

NEW QUESTION 117

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

SAP-C01 Practice Exam Features:

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

100% Actual & Verified — Instant Download, Please Click
[Order The SAP-C01 Practice Test Here](#)