

SAP-C01 Dumps

AWS Certified Solutions Architect- Professional

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



NEW QUESTION 1

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 service
- 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 bucket
- 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 service
- F. Set up an ECS service that is fronted by a cross-zone ALB
- 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 bucket
- 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 function
- M. Store request meta information in DynamoDB with Auto Scaling and static content in a secured S3 bucket
- 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 function
- R. Store request meta information in an Amazon ElastiCache Multi-AZ cluster and static content in a secured S3 bucket
- S. Generate presigned URLs when returning references to content stored in Amazon S3.

Answer: C

NEW QUESTION 2

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

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

Answer: D

Explanation:

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

NEW QUESTION 3

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 4

A company currently uses Amazon EBS and Amazon RDS for storage purposes. The company intends to use a pilot light approach for disaster recovery in a different AWS Region. The company has an RTO of 6 hours and an RPO of 24 hours.

Which solution would achieve the requirements with MINIMAL cost?

- A. Use AWS Lambda to create daily EBS and RDS snapshots, and copy them to the disaster recovery region
- B. Use Amazon Route 53 with active-passive failover configuration
- C. Use Amazon EC2 in an Auto Scaling group with the capacity set to 0 in the disaster recovery region.
- D. Use AWS Lambda to create daily EBS and RDS snapshots, and copy them to the disaster recovery region
- E. Use Amazon Route 53 with active-active failover configuration
- F. Use Amazon EC2 in an AutoScaling group configured in the same way as in the primary region.

- G. Use Amazon ECS to handle long-running tasks to create daily EBS and RDS snapshots, and copy to the disaster recovery regio
- H. Use Amazon Route 53 with active-passive failover configuratio
- I. Use Amazon EC2 in an Auto Scaling group with the capacity set to 0 in the disaster recovery region.
- J. Use EBS and RDS cross-region snapshot copy capability to create snapshots in the disaster recovery regio
- K. Use Amazon Route 53 with active-active failover configuratio
- L. Use Amazon EC2 in an Auto Scaling group with the capacity set to 0 in the disaster recovery region.

Answer: A

Explanation:

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

NEW QUESTION 5

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

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

Which solution is MOST cost-effective?

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

Answer: D

Explanation:

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

NEW QUESTION 6

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 schedules 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 applicatio
- B. Pause the data sources, export the Elasticsearch index from on premises, and import into the EC2 Elasticsearch serve
- 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 applicatio
- E. Use AWS DMS to replicate Elasticsearch dat
- 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 AW
- 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 applicatio
- L. Pause the data source feeds, export the Elasticsearch index from on premises, and import into the Amazon ES cluste
- 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 7

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

Which steps should a Solutions Architect take to reduce costs?

- A. Enable AWS Business Support and review AWS Trusted Advisor's cost check
- B. Create Amazon EC2 Auto Scaling groups for applications that experience fluctuating deman
- C. Save AWS Simple Monthly Calculator reports in Amazon S3 for trend analysi
- D. Create a master account under Organizations and have teams join for consolidating billing.
- E. Enable Cost Explorer and AWS Business Support Reserve Amazon EC2 and Amazon RDS DB instance
- F. Use Amazon CloudWatch and AWS Trusted Advisor for monitoring and to receive cost-savings suggestion
- G. Create a master account under Organizations and have teams join for consolidated billing.
- H. Create an AWS Lambda function that changes the instance size based on Amazon CloudWatch alarms.Reserve instances based on AWS Simple Monthly

Calculator suggestion

- I. Have an AWS Well-Architected framework review and apply recommendation
- J. Create a master account under Organizations and have teams join for consolidated billing.
- K. Create a budget and monitor for costs exceeding the budget
- L. Create Amazon EC2 Auto Scaling groups for applications that experience fluctuating demand
- M. Create an AWS Lambda function that changes instance sizes based on Amazon CloudWatch alarm
- N. Have each team upload their bill to an Amazon S3 bucket for analysis of team spending
- O. Use Spot instances on nightly batch processing jobs.

Answer: D

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

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

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

Answer: A

NEW QUESTION 10

A company would like to implement a serverless application by using Amazon API Gateway, AWS Lambda, and Amazon DynamoDB. They deployed a proof of concept and stated that the average response time is greater than what their upstream services can accept. Amazon CloudWatch metrics did not indicate any issues with DynamoDB but showed that some Lambda functions were hitting their timeout.

Which of the following actions should the Solutions Architect consider to improve performance? (Choose two.)

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

Answer: BD

Explanation:

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

NEW QUESTION 10

A company runs its containerized batch jobs on Amazon ECS. The jobs are scheduled by submitting a container image, a task definition, and the relevant data to an Amazon S3 bucket. Container images may be unique per job. Running the jobs as quickly as possible is of utmost importance, so submitting job artifacts to the S3 bucket triggers the job to run immediately. Sometimes there may be no jobs running at all. However, jobs of any size can be submitted with no prior warning to the IT Operations team. Job definitions include CPU and memory resource requirements.

What solution will allow the batch jobs to complete as quickly as possible after being scheduled?

- A. Schedule the jobs on an Amazon ECS cluster using the Amazon EC2 launch type

- B. Use Service Auto Scaling to increase or decrease the number of running tasks to suit the number of running jobs.
- C. Schedule the jobs directly on EC2 instance
- D. Use Reserved Instances for the baseline minimum load, and use On-Demand Instances in an Auto Scaling group to scale up the platform based on demand.
- E. Schedule the jobs on an Amazon ECS cluster using the Fargate launch type
- F. Use Service Auto Scaling to increase or decrease the number of running tasks to suit the number of running jobs.
- G. Schedule the jobs on an Amazon ECS cluster using the Fargate launch type
- H. Use Spot Instances in an Auto Scaling group to scale the platform based on demand
- I. Use Service Auto Scaling to increase or decrease the number of running tasks to suit the number of running jobs.

Answer: C

NEW QUESTION 11

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 recordings 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 center
- B. Set up an S3 lifecycle policy to move the call recordings to Amazon S3 Glacier after 90 days
- C. Use an AWS Lambda trigger to transcribe the call recordings with Amazon Transcribe
- 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 center
- F. Set up an S3 lifecycle policy to move the call recordings to Amazon S3 Glacier after 90 days
- G. Use an AWS Lambda trigger to transcribe the call recordings with Amazon Mechanical Turk
- 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 application
- J. Upload the call recordings to this application from the call center and store them on an Amazon EFS mount point
- K. Use AWS Backup to archive the call recording after 90 days
- 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 days
- O. Use Amazon EC2 instances in the queue as the scaling metric
- P. Use Amazon S3, Amazon API Gateway, and AWS Lambda to host the review and scoring application.

Answer: B

NEW QUESTION 13

A company has developed a new billing application that will be released in two weeks. Developers are testing the application running on 10 EC2 instances managed by an Auto Scaling group in subnet 172.31.0.0/24 within VPC A with CIDR block 172.31.0.0/16. The Developers noticed connection timeout errors in the application logs while connecting to an Oracle database running on an Amazon EC2 instance in the same region within VPC B with CIDR block 172.50.0.0/16. The IP of the database instance is hard-coded in the application instances. Which recommendations should a Solutions Architect present to the Developers to solve the problem in a secure way with minimal maintenance and overhead?

- A. Disable the SrcDestCheck attribute for all instances running the application and Oracle Database. Change the default route of VPC A to point to the ENI of the Oracle Database that has an IP address assigned within the range of 172.50.0.0/26
- B. Create and attach internet gateways for both VPCs
- C. Configure default routes to the Internet gateways for both VPCs
- D. Assign an Elastic IP for each Amazon EC2 instance in VPC A
- E. Create a VPC peering connection between the two VPCs and add a route to the routing table of VPC A that points to the IP address range of 172.50.0.0/16
- F. Create an additional Amazon EC2 instance for each VPC as a customer gateway; create one virtual private gateway (VGW) for each VPC, configure an end-to-end VPC, and advertise the routes for 172.50.0.0/16

Answer: C

NEW QUESTION 18

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 template. 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 to 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 to 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 23

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

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

Answer: B

NEW QUESTION 24

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 28

A company stores sales transaction data in Amazon DynamoDB tables. To detect anomalous behaviors and respond quickly, all changes to the items stored in the DynamoDB tables must be logged within 30 minutes. Which solution meets the requirements?

- A. Copy the DynamoDB tables into Apache Hive tables on Amazon EMR every hour and analyze them for anomalous behavior
- B. Send Amazon SNS notifications when anomalous behaviors are detected.
- C. Use AWS CloudTrail to capture all the APIs that change the DynamoDB table
- D. Send SNS notifications when anomalous behaviors are detected using CloudTrail event filtering.
- E. Use Amazon DynamoDB Streams to capture and send updates to AWS Lambda
- F. Create a Lambda function to output records to Amazon Kinesis Data Stream
- G. Analyze any anomalies with Amazon Kinesis Data Analytics
- H. Send SNS notifications when anomalous behaviors are detected.
- I. Use event patterns in Amazon CloudWatch Events to capture DynamoDB API call events with an AWS Lambda function as a target to analyze behavior
- J. Send SNS notifications when anomalous behaviors are detected.

Answer: C

Explanation:

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

NEW QUESTION 30

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 mod
- 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 DynamoD
- E. Use an AWS KMS key to encrypt DynamoDB objects at rest.
- F. Deploy instances with Amazon EBS volumes attached to store this dat
- 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 31

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

What is the MOST cost-effective architecture that offers both security and ease of management?

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

Answer: D

Explanation:

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

NEW QUESTION 33

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 storag
- C. Copy the site contents to the volumes at launch and to Amazon S3 at shutdown.
- D. Configure an Amazon CloudFront distributio
- 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 suite contents, and point the distribution at the ALB.

Answer: C

Explanation:

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

NEW QUESTION 36

A company runs an IoT platform on AWS. IoT sensors in various locations send data to the company's Node.js API servers on Amazon EC2 instances running behind an Application Load Balancer. The data is stored in an Amazon RDS MySQL DB instance that uses a 4 TB General Purpose SSD volume.

The number of sensors the company has deployed in the field has increased over time, and is expected to grow significantly. The API servers are consistently overloaded and RDS metrics show high write latency.

Which of the following steps together will resolve the issues permanently and enable growth as new sensors are provisioned, while keeping this platform cost-efficient? (Choose two.)

- A. Resize the MySQL General Purpose SSD storage to 6 TB to improve the volume's IOPS
- B. Re-architect the database tier to use Amazon Aurora instead of an RDS MySQL DB instance and add read replicas
- C. Leverage Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data
- D. Use AWS-X-Ray to analyze and debug application issues and add more API servers to match the load
- E. Re-architect the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance

Answer: CE

NEW QUESTION 37

A company runs a legacy system on a single m4.2xlarge Amazon EC2 instance with Amazon EBS2 storage. The EC2 instance runs both the web server and a self-

managed Oracle database. A snapshot is made of the EBS volume every 12 hours, and an AMI was created from the fully configured EC2 instance. A recent event that terminated the EC2 instance led to several hours of downtime. The application was successfully launched from the AMI, but the age of the EBS snapshot and the repair of the database resulted in the loss of 8 hours of data. The system was also down for 4 hours while the Systems Operators manually performed these processes.

What architectural changes will minimize downtime and reduce the chance of lost data?

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

Answer: B

Explanation:

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

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

NEW QUESTION 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 47

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

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

Which design meets these requirements?

- A. Apply a service control policy (SCP) that allows access to IAM, Amazon RDS, and AWS CloudTrail. Load customer records in Amazon RDS MySQL and train users to execute queries using the AWS CLI

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

Answer: D

NEW QUESTION 48

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 52

A company's application is increasingly popular and experiencing latency because of high volume reads on the database server.

The service has the following properties:

- A highly available REST API hosted in one region using Application Load Balancer (ALB) with auto scaling.
- A MySQL database hosted on an Amazon EC2 instance in a single Availability Zone.

The company wants to reduce latency, increase in-region database read performance, and have multi-region disaster recovery capabilities that can perform a live recovery automatically without any data or performance loss (HA/DR).

Which deployment strategy will meet these requirements?

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

Answer: A

NEW QUESTION 55

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 57

An enterprise company is using a multi-account AWS strategy. There are separate accounts for development, staging, and production workloads. To control costs and improve governance, the following requirements have been defined:

- The company must be able to calculate the AWS costs for each project
- The company must be able to calculate the AWS costs for each environment: development, staging, and production
- Commonly deployed IT services must be centrally managed
- Business units can deploy pre-approved IT services only
- Usage of AWS resources in the development account must be limited

Which combination of actions should be taken to meet these requirements? (Select THREE.)

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

Answer: CEF

NEW QUESTION 60

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 64

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

Which solution meets the company's requirements?

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

Answer: B

Explanation:

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

NEW QUESTION 69

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

Which of the following actions would allow the database to handle the month-end load with the LEAST impact on performance?

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

Answer: B

NEW QUESTION 70

A company's data center is connected to the AWS Cloud over a minimally used 10-Gbps AWS Direct Connect connection with a private virtual interface to its virtual private cloud (VPC). The company internet connection is 200 Mbps and the company has a 150-TB dataset that is created each Friday. The data must be transferred and available in Amazon S3 on Monday morning.

Which is the LEAST expensive way to meet the requirements while allowing for data transfer growth?

- A. Order two 80-GB AWS Snowball appliances. Offload the data to the appliances and ship them to AWS. AWS will copy the data from the Snowball appliances to Amazon S3.
- B. Create a VPC endpoint for Amazon S3. Copy the data to Amazon S3 by using the VPC endpoint, forcing the transfer to use the Direct Connect connection.
- C. Create a VPC endpoint for Amazon S3. Set up a reverse proxy farm behind a Classic Load Balancer in the VPC. Copy the data to Amazon S3 using the proxy.
- D. Create a public virtual interface on a Direct Connect connection and copy the data to Amazon S3 over the connection.

Answer: D

NEW QUESTION 73

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 75

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 80

A Solutions Architect is designing the storage layer for a data warehousing application. The data files are large, but they have statically placed metadata at the beginning of each file that describes the size and placement of the file's index. The data files are read in by a fleet of Amazon EC2 instances that store the index size, index location, and other category information about the data file in a database. That database is used by Amazon EMR to group files together for deeper analysis.

What would be the MOST cost-effective, high availability storage solution for this workflow?

- A. Store the data files in Amazon S3 and use Range GET for each file's metadata, then index the relevant data.
- B. Store the data files in Amazon EFS mounted by the EC2 fleet and EMR nodes.
- C. Store the data files on Amazon EBS volumes and allow the EC2 fleet and EMR to mount and unmount the volumes where they are needed.
- D. Store the content of the data files in Amazon DynamoDB tables with the metadata, index, and data as their own keys.

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/API/RESTObjectGET.html>

NEW QUESTION 81

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

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

Which of the following solutions is MOST cost-effective?

- A. Keep the website on T2 instances. Determine the minimum number of website instances required during off-peak times and use Spot Instances to cover them while using Reserved Instances to cover peak demand. Use Amazon EC2 R4 and Amazon EC2 R5 Reserved Instances in an Auto Scaling group for the video analysis application.
- B. Keep the website on 12 instances. Determine the minimum number of website instances required during off-peak times and use Reserved Instances to cover them while using On-Demand Instances to cover peak demand. Use Spot Fleet for the video analysis application comprised of Amazon EC2 C4 and Amazon EC2 C5 Spot Instances.
- C. Migrate the website to AWS Elastic Beanstalk and Amazon EC2 C4 instances. Determine the minimum number of website instances required during off-peak times and use On-Demand instances to cover them while using Spot capacity to cover peak demand. Use Spot Fleet for the video analysis application comprised of C4 and Amazon EC2 C5 instances.
- D. Migrate the website to AWS Elastic Beanstalk and Amazon EC2 R4 instances. Determine the minimum number of website instances required during off-peak times and use Reserved Instances to cover them while using On-Demand Instances to cover peak demand. Use Spot Fleet for the video analysis application comprised of R4 and Amazon EC2 R5 instances.

Answer: B

NEW QUESTION 84

A company plans to move regulated and security-sensitive businesses to AWS. The Security team is developing a framework to validate the adoption of AWS best practice and industry-recognized compliance standards. The AWS Management Console is the preferred method for teams to provision resources.

Which strategies should a Solutions Architect use to meet the business requirements and continuously assess, audit, and monitor the configurations of AWS resources? (Choose two.)

- A. Use AWS Config rules to periodically audit changes to AWS resources and monitor the compliance of the configuration.
- B. Develop AWS Config custom rules using AWS Lambda to establish a test-driven development approach, and further automate the evaluation of configuration changes against the required controls.
- C. Use Amazon CloudWatch Logs agent to collect all the AWS SDK log.
- D. Search the log data using a pre-defined set of filter patterns that machines mutating API call.
- E. Send notifications using Amazon CloudWatch alarms when unintended changes are performed.
- F. Archive log data by using a batch export to Amazon S3 and then Amazon Glacier for a long-term retention and auditability.
- G. Use AWS CloudTrail events to assess management activities of all AWS account.
- H. Ensure that CloudTrail is enabled in all accounts and available AWS service.
- I. Enable trails, encrypt CloudTrail event log files with an AWS KMS key, and monitor recorded activities with CloudWatch Logs.
- J. Use the Amazon CloudWatch Events near-real-time capabilities to monitor system events patterns, and trigger AWS Lambda functions to automatically revert non-authorized changes in AWS resource.
- K. Also, target Amazon SNS topics to enable notifications and improve the response time of incident responses.
- L. Use CloudTrail integration with Amazon SNS to automatically notify unauthorized API activities. Ensure that CloudTrail is enabled in all accounts and available AWS service.
- M. Evaluate the usage of Lambda functions to automatically revert non-authorized changes in AWS resources.

Answer: AC

Explanation:

<https://docs.aws.amazon.com/awsccloudtrail/latest/userguide/cloudwatch-alarms-for-cloudtrail.html>

https://docs.aws.amazon.com/en_pv/awscloudtrail/latest/userguide/best-practices-security.html

NEW QUESTION 89

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

- A. Define AWS IAM roles based on the functional responsibilities of the users in a central account
- B. Create a SAML-based identity management provider
- C. Map users in the on-premises groups to IAM role
- D. Establish trust relationships between the other accounts and the central account.
- E. Deploy a common set of AWS IAM users, groups, roles, and policies in all of the AWS accounts using AWS Organization
- F. Implement federation between the on-premises identity provider and the AWS accounts.
- G. Use AWS Organizations in a centralized account to define service control policies (SCPs). Create a SAML-based identity management provider in each account and map users in the on-premises groups to AWS IAM roles.
- H. Perform a thorough analysis of the user base and create AWS IAM users accounts that have the necessary 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 91

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 branch
- 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 branch
- 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 pipeline
- 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 96

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 98

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

- Consolidate all accounts into one organization.
- Allow full access to the Amazon EC2 service from the master account and the secondary accounts.
- Minimize the effort required to add additional secondary accounts.

Which combination of steps should be included in the solution? (Choose two.)

- A. Create an organization from the master account
- B. Send invitations to the secondary accounts from the master account
- C. Accept the invitations and create an OU.
- D. Create an organization from the master account

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

Answer: AD

Explanation:

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

NEW QUESTION 103

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

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

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

- A. Create accounts for each production workload within an organization in AWS Organizations Place the production accounts within an organizational unit (OU) For each account delete the default VPC Create an SCP with a Deny rule for the attach an internet gateway and create a default VPC actions Attach the SCP to the OU for the production accounts
- B. Create accounts for each production workload within an organization in AWS Organizations Place the production accounts within an organizational unit (OU) Create an SCP with a Deny rule on the attach an internet gateway action Create an SCP with a Deny rule to prevent use of the default VPC Attach the SCPs to the OU for the production accounts
- C. Create a SCP containing a Deny Effect for cloudfront". lam:*, route53* and support* with a StringNotEquals condition on an aws RequestedRegion condition key with us-west-2 and eu-central-1 values Attach the SCP to the organization's root.
- D. Create an IAM permission boundary containing a Deny Effect for cloudfront'. lam * route53' and support" with a StringNotEquals condition on an aws RequestedRegion condition key with us-west 2 and eu-central-1 values Attach the permission boundary to an IAM group containing the development and production users.
- E. Create accounts for each development workload within an organization 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 106

A company is running a high-user-volume media-sharing application on premises It currently hosts about 400 TB of data with millions of video files The company is migrating this application to AWS to improve reliability and reduce costs

The Solutions Architecture team plans to store the videos in an Amazon S3 bucket and use Amazon

CloudFront to distribute videos to users. The company needs to migrate this application to AWS within 10 days with the least amount of downtime possible. The company currently has 1 Gbps connectivity to the internet with 30 percent free capacity

Which of the following solutions would enable the company to migrate the workload to AWS and meet an of the requirements?

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

Answer: D

Explanation:

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

NEW QUESTION 107

A company has implemented AWS Organizations. It has recently set up a number of new accounts and wants to deny access to a specific set of AWS services in these new accounts.

How can this be controlled MOST efficiently?

- A. Create an IAM policy in each account that denies access to the service
- B. Associate the policy with an IAM group, and add all IAM users to the group.
- C. Create a service control policy that denies access to the service
- D. Add all of the new accounts to a single organizations unit (OU), and apply the policy to that OU.
- E. Create an IAM policy in each account that denies access to the service
- F. Associate the policy with an IAM role, and instruct users to log in using their corporate credentials and assume the IAM role.
- G. Create a service control policy that denies access to the services, and apply the policy to the root of the organization.

Answer: B

NEW QUESTION 110

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

What is the MOST secure deployment design that meets all solution requirements?

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

Answer: A

NEW QUESTION 114

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

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

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

Answer: B

NEW QUESTION 116

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 117

A company is planning the migration of several lab environments used for software testing. An assortment of custom tooling is used to manage the test runs for each lab. The labs use immutable infrastructure for the software test runs, and the results are stored in a highly available SQL database cluster. Although completely rewriting the custom tooling is out of scope for the migration project, the company would like to optimize workloads during the migration.

Which application migration strategy meets this requirement?

- A. Re-host
- B. Re-platform
- C. Re-factor/re-architect
- D. Retire

Answer: B

Explanation:

<https://aws.amazon.com/blogs/enterprise-strategy/6-strategies-for-migrating-applications-to-the-cloud/>

NEW QUESTION 121

The Security team needs to provide a team of interns with an AWS environment so they can build the serverless video transcoding application. The project will use Amazon S3, AWS Lambda, Amazon API Gateway, Amazon Cognito, Amazon DynamoDB, and Amazon Elastic Transcoder.

The interns should be able to create and configure the necessary resources, but they may not have access to create or modify AWS IAM roles. The Solutions Architect creates a policy and attaches it to the interns' group.

How should the Security team configure the environment to ensure that the interns are self-sufficient?

- A. Create a policy that allows creation of project-related resources only

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

Answer: A

NEW QUESTION 124

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 127

A company has a data center that must be migrated to AWS as quickly as possible. The data center has a 500 Mbps AWS Direct Connect link and a separate, fully available 1 Gbps ISP connection. A Solutions Architect must transfer 20 TB of data from the data center to an Amazon S3 bucket.

What is the FASTEST way transfer the data?

- A. Upload the data to the S3 bucket using the existing DX link.
- B. Send the data to AWS using the AWS Import/Export service.
- C. Upload the data using an 80 TB AWS Snowball device.
- D. Upload the data to the S3 bucket using S3 Transfer Acceleration.

Answer: D

Explanation:

<https://aws.amazon.com/s3/faqs/>

NEW QUESTION 130

A company is having issues with a newly deployed server less infrastructure that uses Amazon API Gateway, Amazon Lambda, and Amazon DynamoDB.

In a steady state, the application performs as expected. However, during peak load, tens of thousands of simultaneous invocations are needed and user requests fail multiple times before succeeding. The company has checked the logs for each component, focusing specifically on Amazon CloudWatch Logs for Lambda. There are no errors logged by the services or applications.

What might cause this problem?

- A. Lambda has very memory assigned, which causes the function to fail at peak load.
- B. Lambda is in a subnet that uses a NAT gateway to reach out to the internet, and the function instance does not have sufficient Amazon EC2 resources in the VPC to scale with the load.
- C. The throttle limit set on API Gateway is very low during peak load, the additional requests are not making their way through to Lambda.
- D. DynamoDB is set up in an auto scaling mode.
- E. During peak load, DynamoDB adjusts capacity and through successfully.

Answer: A

NEW QUESTION 135

A company has a single AWS master billing account, which is the root of the AWS Organizations hierarchy. The company has multiple AWS accounts within this hierarchy, all organized into organization units (OUs). More OUs and AWS accounts will continue to be created as other parts of the business migrate applications to AWS. These business units may need to use different AWS services. The Security team is implementing the following requirements for all current and future AWS accounts.

* Control policies must be applied across all accounts to prohibit AWS servers.

* Exceptions to the control policies are allowed based on valid use cases. Which solution will meet these requirements with minimal optional overhead?

- A. Use an SCP in Organizations to implement a deny list of AWS server
- B. Apply this SCP at the level
- C. For any specific exceptions for an OU, create a new SCP for that OU and add the required AWS services to the allow list.
- D. Use an SCP in Organizations to implement a deny list of AWS service
- E. Apply this SCP at the root level and each OU
- F. Remove the default AWS managed SCP from the root level and all OU level
- G. For any specific exceptions, modify the SCP attached to that OU, and add the required AWS services to the allow list.
- H. Use an SCP in Organization to implement a deny list of AWS service
- I. Apply this SCP at each OU level
- J. Leave the default AWS managed SCP at the root level. For any specific exceptions for an OU, create a new SCP for that OU.
- K. Use an SCP in Organizations to implement an allow list of AWS service
- L. Apply this SCP at the root level
- M. Remove the default AWS managed SCP from the root level and all OU level
- N. For any specific exceptions for an OU, modify the SCP attached to that OU, and add the required AWS services to the allow list.

Answer: B

NEW QUESTION 137

A company that is new to AWS reports it has exhausted its service limits across several accounts that are on the Basic Support plan. The company would like to prevent this from happening in the future.

What is the MOST efficient way of monitoring and managing all service limits in the company's accounts?

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

Answer: D

Explanation:

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

NEW QUESTION 141

A company is running multiple applications on Amazon EC2. Each application is deployed and managed by multiple business units. All applications are deployed on a single AWS account but on different virtual private clouds (VPCs). The company uses a separate VPC in the same account for test and development purposes.

Production applications suffered multiple outages when users accidentally terminated and modified resources that belonged to another business unit. A Solutions Architect has been asked to improve the availability of the company applications while allowing the Developers access to the resources they need.

Which option meets the requirements with the LEAST disruption?

- A. Create an AWS account for each business 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 145

A Solutions Architect is designing a highly available and reliable solution for a cluster of Amazon EC2 instances.

The Solutions Architect must ensure that any EC2 instance within the cluster recovers automatically after a system failure. The solution must ensure that the recovered instance maintains the same IP address.

How can these requirements be met?

- A. Create an AWS Lambda script to restart any EC2 instances that shut down unexpectedly.
- B. Create an Auto Scaling group for each EC2 instance that has a minimum and maximum size of 1.
- C. Create a new t2.micro instance to monitor the cluster instance
- D. Configure the t2.micro instance to issue an `aws ec2 reboot-instances` command upon failure.
- E. Create an Amazon CloudWatch alarm for the `StatusCheckFailed_System` metric, and then configure an EC2 action to recover the instance.

Answer: B

Explanation:

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

NEW QUESTION 149

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 mod
- 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 day
- 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 151

A company has developed a web application that runs on Amazon EC2 instances in one AWS Region. The company has taken on new business in other countries and must deploy its application into other to meet low-latency requirements for its users. The regions can be segregated, and an application running in one region does not need to communicate with instances in other regions.

How should the company's Solutions Architect automate the deployment of the application so that it can be MOST efficiently deployed into multiple regions?

- A. Write a bash script that uses the AWS CLI to query the current state in one region and output a JSON representatio
- B. Pass the JSON representation to the AWS CLI, specifying the --region parameter to deploy the application to other regions.
- C. Write a bash script that uses the AWS CLI to query the current state in one region and output an AWS CloudFormation templat
- D. Create a CloudFormation stack from the template by using the AWS CLI, specifying the --region parameter to deploy the application to other regions.
- E. Write a CloudFormation template describing the application's infrastructure in the resources section. Create a CloudFormation stack from the template by using the AWS CLI, specify multiple regions using the --regions parameter to deploy the application.
- F. Write a CloudFormation template describing the application's infrastructure in the Resources section. Use a CloudFormation stack set from an administrator account to launch stack instances that deploy the application to other regions.

Answer: D

Explanation:

A stack set lets you create stacks in AWS accounts across regions by using a single AWS CloudFormation template. All the resources included in each stack are defined by the stack set's AWS CloudFormation template. As you create the stack set, you specify the template to use, as well as any parameters and capabilities that template requires. <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-concepts.html>
<https://sanderknape.com/2017/07/cloudformation-stacksets-automated-cross-account-region-deployments/>

NEW QUESTION 155

A financial company is using a high-performance compute cluster running on Amazon EC2 instances to perform market simulations A DNS record must be created in an Amazon Route 53 private hosted zone when instances start The DNS record must be removed after instances are terminated.

Currently the company uses a combination of Amazon CloudWatch Events and AWS Lambda to create the DNS record. The solution worked well in testing with small clusters, but in production with clusters containing thousands of instances the company sees the following error in the Lambda logs:

HTTP 400 error (Bad request).

The response header also includes a status code element with a value of "Throttling" and a status message element with a value of "Rate exceeded "

Which combination of steps should the Solutions Architect take to resolve these issues? (Select THREE)

- A. Configure an Amazon SQS FIFO queue and configure a CloudWatch Events rule to use this queue as a target
- B. Remove the Lambda target from the CloudWatch Events rule
- C. Configure an Amazon Kinesis data stream and configure a CloudWatch Events rule to use this queue as a target Remove the Lambda target from the CloudWatch Events rule
- D. Update the CloudWatch Events rule to trigger on Amazon EC2 "Instance Launch Successful" and "Instance Terminate Successful" events for the Auto Scaling group used by the cluster
- E. Configure a Lambda function to retrieve messages from an Amazon SQS queue Modify the Lambda function to retrieve a maximum of 10 messages then batch the messages by Amazon Route 53 API call type and submit Delete the messages from the SQS queue after successful API calls.
- F. Configure an Amazon SQS standard queue and configure the existing CloudWatch Events rule to use this queue as a target Remove the Lambda target from the CloudWatch Events rule.
- G. Configure a Lambda function to read data from the Amazon Kinesis data stream and configure the batch window to 5 minutes Modify the function to make a single API call to Amazon Route 53 with all records read from the kinesis data stream

Answer: BEF

NEW QUESTION 160

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your SAP-C01 Exam with Our Prep Materials Via below:

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