

## SAP-C01 Dumps

### AWS Certified Solutions Architect- Professional

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



**NEW QUESTION 1**

A company is currently running a production workload on AWS that is very I/O intensive. Its workload consists of a single tier with 10 c4.8xlarge instances, each with 2 TB gp2 volumes. The number of processing jobs has recently increased, and latency has increased as well. The team realizes that they are constrained on the IOPS. For the application to perform efficiently, they need to increase the IOPS by 3,000 for each of the instances. Which of the following designs will meet the performance goal MOST cost effectively?

- A. Change the type of Amazon EBS volume from gp2 to io1 and set provisioned IOPS to 9,000.
- B. Increase the size of the gp2 volumes in each instance to 3 TB.
- C. Create a new Amazon EFS file system and move all the data to this new file system
- D. Mount this file system to all 10 instances.
- E. Create a new Amazon S3 bucket and move all the data to this new bucket
- F. Allow each instance to access this S3 bucket and use it for storage.

**Answer: B**

**NEW QUESTION 2**

A Solutions Architect must design a highly available, stateless, REST service. The service will require multiple persistent storage layers for service object meta information and the delivery of content. Each request needs to be authenticated and securely processed. There is a requirement to keep costs as low as possible? How can these requirements be met?

- A. Use AWS Fargate to host a container that runs a self-contained REST 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 API
- 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 API
- 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 API
- 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 API
- 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 3**

A company receives clickstream data files to Amazon S3 every five minutes. A Python script runs as a cron job once a day on an Amazon EC2 instance to process each file and load it into a database hosted on Amazon RDS. The cron job takes 15 to 30 minutes to process 24 hours of data. The data consumers ask for the data to be available as soon as possible. Which solution would accomplish the desired outcome?

- A. Increase the size of the instance to speed up processing and update the schedule to run once an hour.
- B. Convert the cron job to an AWS Lambda function and trigger this new function using a cron job on an EC2 instance.
- C. Convert the cron job to an AWS Lambda function and schedule it to run once an hour using Amazon CloudWatch events.
- D. Create an AWS Lambda function that runs when a file is delivered to Amazon S3 using S3 event notifications.

**Answer: D**

**Explanation:**

<https://docs.aws.amazon.com/lambda/latest/dg/with-s3.html>

**NEW QUESTION 4**

A retail company processes point-of-state data on application servers in its data center and writes outputs to Amazon DynamoDB table. The data center is connected to the company's VPC with an AWS Direct Connect (DX) connection, and the application servers require a consistent network connection at speed greater than 2 Gbps.

The company decides that the DynamoDB table needs to be highly available and fault tolerant. The company policy states that the data should be available across two regions.

What changes should the company make to meet these requirements?

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

**Answer:** A

#### NEW QUESTION 5

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

- An application tier that contains 8 m4.xlarge instances
- A Classic Load Balancer
- Amazon S3 as a persistent data store

After one of the EC2 instances fails, users report very slow processing of their requests. A Solutions Architect must recommend design changes to maximize system reliability. The solution must minimize costs.

What should the Solution Architect recommend?

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

**Answer:** B

#### Explanation:

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

#### NEW QUESTION 6

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

- The website should be responsive.
- The website should offer minimal latency.
- The website should be highly available.
- Users should be able to authenticate through social identity providers such as Google, Facebook, and Amazon.
- There should be baseline DDoS protections for spikes in traffic.

How can the design requirements be met?

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

**Answer:** C

#### NEW QUESTION 7

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 8

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 region
- H. Use Amazon Route 53 with active-passive failover configuration
- 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 region
- K. Use Amazon Route 53 with active-active failover configuration
- 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](https://docs.aws.amazon.com/AmazonECS/latest/developerguide/scheduling_tasks.html)

**NEW QUESTION 9**

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

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

**Answer:** D

**NEW QUESTION 10**

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 rebuilt 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 demand
- C. Save AWS Simple Monthly Calculator reports in Amazon S3 for trend analysis
- D. Create a master account under Organizations and have teams join for consolidated 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 suggestions
- 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 recommendations
- 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 10**

While debugging a backend application for an IoT system that supports globally distributed devices a Solutions Architect notices that stale data is occasionally being sent to user devices. Devices often share data, and stale data does not cause issues in most cases. However, device operations are disrupted when a device reads the stale data after an update.

The global system has multiple identical application stacks deployed in different AWS Regions. If a user device travels out of its home geographic region it will always connect to the geographically closest AWS Region to write or read data. The same data is available in all supported AWS Regions using an Amazon DynamoDB global table.

What change should be made to avoid causing disruptions in device operations?

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

E. Select one AWS Region as a master and perform all writes in that AWS Region only Update the backend to use strongly consistent reads

**Answer: B**

#### NEW QUESTION 14

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

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

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

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

How should the Solutions Architect design the solution?

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

**Answer: C**

#### NEW QUESTION 16

A company CFO recently analyzed the company's AWS monthly bill and identified an opportunity to reduce the cost for AWS Elastic Beanstalk environments in use. The CFO has asked a Solutions Architect to design a highly available solution that will spin up an Elastic Beanstalk environment in the morning and terminate it at the end of the day.

The solution should be designed with minimal operational overhead and to minimize costs. It should also be able to handle the increased use of Elastic Beanstalk environments among different teams, and must provide a one-stop scheduler solution for all teams to keep the operational costs low.

What design will meet these requirements?

- A. Set up a Linux EC2 Micro instance
- B. Configure an IAM role to allow the start and stop of the Elastic Beanstalk environment and attach it to the instance
- C. Create scripts on the instance to start and stop the Elastic Beanstalk environment
- D. Configure cron jobs on the instance to execute the scripts.
- E. Develop AWS Lambda functions to start and stop the Elastic Beanstalk environment
- F. Configure a Lambda execution role granting Elastic Beanstalk environment start/stop permissions, and assign the role to the Lambda function
- G. Configure cron expression Amazon CloudWatch Events rules to trigger the Lambda functions.
- H. Develop an AWS Step Functions state machine with "wait" as its type to control the start and stop time. Use the activity task to start and stop the Elastic Beanstalk environment
- I. Create a role for Step Functions to allow it to start and stop the Elastic Beanstalk environment
- J. Invoke Step Functions daily.
- K. Configure a time-based Auto Scaling group
- L. In the morning, have the Auto Scaling group scale up an Amazon EC2 instance and put the Elastic Beanstalk environment start command in the EC2 instance user data
- M. At the end of the day, scale down the instance number to 0 to terminate the EC2 instance.

**Answer: B**

#### Explanation:

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

#### NEW QUESTION 19

A company is migrating its marketing website and content management system from an on-premises data center to AWS. The company wants the AWS application to be developed in a VPC with Amazon EC2 instances used for the web servers and an Amazon RDS instance for the database.

The company has a runbook document that describes the installation process of the on-premises system. The company would like to base the AWS system on the processes referenced in the runbook document. The runbook document describes the installation and configuration of the operating systems, network settings, the website, and content management system software on the servers. After the migration is complete, the company wants to be able to make changes quickly to take advantage of other AWS features.

How can the application and environment be deployed and automated in AWS, while allowing for future changes?

- A. Update the runbook to describe how to create the VPC, the EC2 instances, and the RDS instance for the application by using the AWS Console
- B. Make sure that the rest of the steps in the runbook are updated to reflect any changes that may come from the AWS migration.

- C. Write a Python script that uses the AWS API to create the VPC, the EC2 instances, and the RDS instance for the applicatio
- D. Write shell scripts that implement the rest of the steps in the runboo
- E. Have the Python script copy and run the shell scripts on the newly created instances to complete the installation.
- F. Write an AWS CloudFormation template that creates the VPC, the EC2 instances, and the RDS instance for the applicatio
- G. Ensure that the rest of the steps in the runbook are updated to reflect any changes that may come from the AWS migration.
- H. Write an AWS CloudFormation template that creates the VPC, the EC2 instances, and the RDS instance for the applicatio
- I. Include EC2 user data in the AWS CloudFormation template to install and configure the software.

**Answer:** D

**NEW QUESTION 23**

A large company has increased its utilization of AWS over time in an unmanaged way. As such, they have a large number of independent AWS accounts across different business units, projects, and environments. The company has created a Cloud Center of Excellence team, which is responsible for managing all aspects of the AWS Cloud, including their AWS accounts.

Which of the following should the Cloud Center of Excellence team do to BEST address their requirements in a centralized way? (Select two.)

- A. Control all AWS account root user credential
- B. Assign AWS IAM users in the account of each user who needs to access AWS resource
- C. Follow the policy of least privilege in assigning permissions to each user.
- D. Tag all AWS resources with details about the business unit, project, and environmen
- E. Send all AWS Cost and Usage reports to a central Amazon S3 bucket, and use tools such as Amazon Athena and Amazon QuickSight to collect billing details by business unit.
- F. Use the AWS Marketplace to choose and deploy a Cost Management too
- G. Tag all AWS resources with details about the business unit, project, and environmen
- H. Send all AWS Cost and Usage reports for the AWS accounts to this tool for analysis.
- I. Set up AWS Organization
- J. Enable consolidated billing, and link all existing AWS accounts to a master billing accoun
- K. Tag all AWS resources with details about the business unit, project and environmen
- L. Analyze Cost and Usage reports using tools such as Amazon Athena and Amazon QuickSight to collect billing details by business unit.
- M. Using a master AWS account, create IAM users within the master accoun
- N. Define IAM roles in the other AWS accounts, which cover each of the required functions in the accoun
- O. Follow the policy of least privilege in assigning permissions to each role, then enable the IAM users to assume the roles that they need to use.

**Answer:** DE

**NEW QUESTION 28**

A company will several AWS accounts is using AWS Organizations and service control policies (SCPs). An Administrator created the following SCP and has attached it to an organizational unit (OU) that contains AWS account 1111-1111-1111:

```
{
  "Version": "2012-10-27",
  "Statement": [
    {
      "Sid": "AllowsAllActions",
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*"
    },
    {
      "Sid": "DenyCloudTrail",
      "Effect": "Deny",
      "Action": "cloudtrail:*",
      "Resource": "*"
    }
  ]
}
```

Developers working in account 1111-1111-1111 complain that they cannot create Amazon S3 buckets. How should the Administrator address this problem?

- A. Add s3:CreateBucket with "Allow" effect to the SCP.
- B. Remove the account from the OU, and attach the SCP directly to account 1111-1111-1111.
- C. Instruct the Developers to add Amazon S3 permissions to their IAM entities.
- D. Remove the SCP from account 1111-1111-1111.

**Answer:** C

**NEW QUESTION 29**

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 cente

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

**Answer: B**

#### NEW QUESTION 32

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

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

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

**Answer: AC**

#### NEW QUESTION 34

A company runs a public-facing application that uses a Java-based web service via a RESTful API. It is hosted on Apache Tomcat on a single server in a data center that runs consistently at 30% CPU utilization. Use of the API is expected to increase by 10 times with a new product launch. The business wants to migrate the application to AWS with no disruption and needs it to scale to meet demand.

The company has already decided to use Amazon Route 53 and CNAME records to redirect traffic. How can these requirements be met with the LEAST amount of effort?

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

**Answer: A**

#### NEW QUESTION 36

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

How can this be accomplished?

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

**Answer: B**

#### Explanation:

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

#### NEW QUESTION 37

A company runs a three-tier application in AWS. Users report that the application performance can vary greatly depending on the time of day and functionality being accessed.

The application includes the following components:

- Eight t2.large front-end web servers that serve static content and proxy dynamic content from the application tier.
- Four t2.large application servers.
- One db.m4.large Amazon RDS MySQL Multi-AZ DB instance.

Operations has determined that the web and application tiers are network constrained.

Which of the following should cost effective improve application performance? (Choose two.)

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

**Answer:** BD

**Explanation:**

<https://aws.amazon.com/ec2/instance-types/>

#### NEW QUESTION 41

A large multinational company runs a timesheet application on AWS that is used by staff across the world. The application runs on Amazon EC2 instances in an Auto Scaling group behind an Elastic Load Balancing (ELB) load balancer, and stores in an Amazon RDS MySQL Multi-AZ database instance.

The CFO is concerned about the impact on the business if the application is not available. The application must not be down for more than two hours, but the solution must be as cost-effective as possible.

How should the Solutions Architect meet the CFO's requirements while minimizing data loss?

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

**Answer:** D

#### NEW QUESTION 43

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 47

A company wants to ensure that the workloads for each of its business units have complete autonomy and a minimal blast radius in AWS. The Security team must be able to control access to the resources and services in the account to ensure that particular services are not used by the business units.

How can a Solutions Architect achieve the isolation requirements?

- A. Create individual accounts for each business unit and add the account to an OU in AWS Organizations. Modify the OU to ensure that the particular services are blocked
- B. Federate each account with an IdP, and create separate roles for the business units and the Security team.
- C. Create individual accounts for each business unit
- D. Federate each account with an IdP and create separate roles and policies for business units and the Security team.
- E. Create one shared account for the entire company
- F. Create separate VPCs for each business unit
- G. Create individual IAM policies and resource tags for each business unit
- H. Federate each account with an IdP, and create separate roles for the business units and the Security team.

- I. Create one shared account for the entire company
- J. Create individual IAM policies and resource tags for each business unit
- K. Federate the account with an IdP, and create separate roles for the business units and the Security team.

**Answer:** A

#### NEW QUESTION 50

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

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

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

**Answer:** B

#### Explanation:

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

#### NEW QUESTION 54

A company has an application that uses Amazon EC2 instances in an Auto Scaling group. The Quality Assurance (QA) department needs to launch a large number of short-lived environments to test the application. The application environments are currently launched by the Manager of the department using an AWS CloudFormation template. To launch the stack, the Manager uses a role with permission to use CloudFormation, EC2 and Auto Scaling APIs. The Manager wants to allow testers to launch their own environments, but does not want to grant broad permission to each user. Which set up would achieve these goals?

- A. Upload the AWS CloudFormation template to Amazon S3. Give users in the QA department permission to assume the Manager's role and add a policy that restricts the permissions to the template and the resources it creates
- B. Train users to launch the template from the CloudFormation console.
- C. Create an AWS Service Catalog product from the environment template
- D. Add a launch constraint to the product with the existing role
- E. Give users in the QA department permission to use AWS Service Catalog APIs only
- F. Train users to launch the templates from the AWS Service Catalog console.
- G. Upload the AWS CloudFormation template to Amazon S3. Give users in the QA department permission to use CloudFormation and S3 APIs, with conditions that restrict the permission to the template and the resources it creates
- H. Train users to launch the template from the CloudFormation console.
- I. Create an AWS Elastic Beanstalk application from the environment template
- J. Give users in the QA department permission to use Elastic Beanstalk permissions only
- K. Train users to launch Elastic Beanstalk environments with the Elastic Beanstalk CLI, passing the existing role to the environment as a service role.

**Answer:** B

#### Explanation:

<https://aws.amazon.com/blogs/mt/how-to-launch-secure-and-governed-aws-resources-with-aws-cloudformation->

#### NEW QUESTION 59

A retail company has a custom .NET web application running on AWS that uses Microsoft SQL Server for the database. The application servers maintain a user's session locally.

Which combination of architecture changes are needed to ensure all tiers of the solution are highly available? (Select THREE.)

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

**Answer:** BDE

#### NEW QUESTION 60

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 sharing

- 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 sharing.
- 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 65**

An online e-commerce business is running a workload on AWS. The application architecture includes a web tier, an application tier for business logic, and a database tier for user and transactional data management. The database server has a 100 GB memory requirement. The business requires cost-efficient disaster recovery for the application with an RTO of 5 minutes and an RPO of 1 hour. The business also has a regulatory requirement for out-of-region disaster recovery with a minimum distance between the primary and alternate sites of 250 miles.

Which of the following options can the Solutions Architect design to create a comprehensive solution for this customer that meets the disaster recovery requirements?

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

**Answer:** C

**NEW QUESTION 66**

A company deployed a three-tier web application in two regions: us-east-1 and eu-west-1. The application must be active in both regions at the same time. The database tier of the application uses a single Amazon RDS Aurora database globally, with a master in us-east-1 and a read replica in eu-west-1. Both regions are connected by a VPN.

The company wants to ensure that the application remains available even in the event of a region-level failure of all of the application's components. It is acceptable for the application to be in read-only mode for up to 1 hour. The company plans to configure two Amazon Route 53 record sets, one for each of the regions.

How should the company complete the configuration to meet its requirements while providing the lowest latency for the application end-users? (Choose two.)

- A. Use failover routing and configure the us-east-1 record set as primary and the eu-west-1 record set as secondary.
- B. Configure an HTTP health check for the web application in us-east-1, and associate it to the us-east-1 record set.
- C. Use weighted routing and configure each record set with a weight of 50. Configure an HTTP health check for each region, and attach it to the record set for that region.
- D. Use latency-based routing for both record sets.
- E. Configure a health check for each region and attach it to the record set for that region.
- F. Configure an Amazon CloudWatch alarm for the health checks in us-east-1, and have it invoke an AWS Lambda function that promotes the read replica in eu-west-1.
- G. Configure an Amazon RDS event notification to react to the failure of the database in us-east-1 by invoking an AWS Lambda function that promotes the read replica in eu-west-1.

**Answer:** CE

**Explanation:**

<https://docs.aws.amazon.com/lambda/latest/dg/services-rds.html>

**NEW QUESTION 70**

An advisory firm is creating a secure data analytics solution for its regulated financial services users. Users will upload their raw data to an Amazon S3 bucket, where they have PutObject permissions only. Data will be analyzed by applications running on an Amazon EMR cluster launched in a VPC. The firm requires that the environment be isolated from the internet. All data at rest must be encrypted using keys controlled by the firm.

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

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

**Answer:** AC

**NEW QUESTION 72**

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 73**

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 76**

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 79**

An enterprise runs 103 line-of-business applications on virtual machines in an on-premises data center. Many of the applications are simple PHP, Java, or Ruby web applications, are no longer actively developed, and serve little traffic.

Which approach should be used to migrate these applications to AWS with the LOWEST infrastructure costs ?

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 80**

A company is creating an account strategy so that they can begin using AWS. The Security team will provide each team with the permissions they need to follow the principle of least privileged access. Teams would like to keep their resources isolated from other groups, and the Finance team would like each team's resource usage separated for billing purposes.

Which account creation process meets these requirements and allows for changes?

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

**Answer: B**

**NEW QUESTION 83**

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

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

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

**Answer: B**

**NEW QUESTION 88**

A company runs a video processing platform. Files are uploaded by users who connect to a web server, which stores them on an Amazon EFS share. This web server is running on a single Amazon EC2 instance. A different group of instances, running in an Auto Scaling group, scans the EFS share directory structure for new files to process and generates new videos (thumbnails, different resolution, compression, etc.) according to the instructions file, which is uploaded along with the video files. A different application running on a group of instances managed by an Auto Scaling group processes the video files and then deletes them from the EFS share. The results are stored in an S3 bucket. Links to the processed video files are emailed to the customer.

The company has recently discovered that as they add more instances to the Auto Scaling Group, many files are processed twice, so image processing speed is not improved. The maximum size of these video files is 2GB.

What should the Solutions Architect do to improve reliability and reduce the redundant processing of video files?

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

**Answer: A**

**NEW QUESTION 92**

A company has a High Performance Computing (HPC) cluster in its on-premises data center which runs thousands of jobs in parallel for one week every month, processing petabytes of images. The images are stored on a network file server, which is replicated to a disaster recovery site. The on-premises data center has reached capacity and has started to spread the jobs out over the course of month in order to better utilize the cluster, causing a delay in the job completion.

The company has asked its Solutions Architect to design a cost-effective solution on AWS to scale beyond the current capacity of 5,000 cores and 10 petabytes of data. The solution must require the least amount of management overhead and maintain the current level of durability.

Which solution will meet the company's requirements?

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

**Answer: B**

#### NEW QUESTION 95

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

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

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

**Answer: C**

#### Explanation:

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

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

[https://docs.aws.amazon.com/dms/latest/userguide/CHAP\\_Security.html](https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Security.html)

#### NEW QUESTION 96

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

What is the LEAST complex method of migrating the database securely and reliably?

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

**Answer: B**

#### NEW QUESTION 97

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 102**

A company is refactoring an existing web service that provides read and write access to structured data. The service must respond to short but significant spikes in the system load. The service must be fault tolerant across multiple AWS Regions.

Which actions should be taken to meet these requirements?

- A. Store the data in Amazon DocumentDB. Create a single global Amazon CloudFront distribution with a custom origin built on edge-optimized Amazon API Gateway and AWS Lambda. Assign the company's domain as an alternate domain for the distribution.
- B. and configure Amazon Route 53 with an alias to the CloudFront distribution.
- C. Store the data in replicated Amazon S3 buckets in two Regions. Create an Amazon CloudFront distribution in each Region, with custom origins built on Amazon API Gateway and AWS Lambda launched in each Region. Assign the company's domain as an alternate domain for both distributions and configure Amazon Route 53 with a failover routing policy between them.
- D. Store the data in an Amazon DynamoDB global table in two Regions using on-demand capacity mode. In both Regions, run the web service as Amazon ECS Fargate tasks in an Auto Scaling ECS service behind an Application Load Balancer (ALB). In Amazon Route 53, configure an alias record in the company's domain and a Route 53 latency-based routing policy with health checks to distribute traffic between the two ALBs.

**Answer:** A

**NEW QUESTION 104**

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 107**

A company is migrating to the cloud. It wants to evaluate the configurations of virtual machines in its existing data center environment to ensure that it can size new Amazon EC2 instances accurately. The company wants to collect metrics, such as CPU, memory, and disk utilization, and it needs an inventory of what processes are running on each instance. The company would also like to monitor network connections to map communications between servers.

Which would enable the collection of this data MOST cost effectively?

- A. Use AWS Application Discovery Service and deploy the data collection agent to each virtual machine in the data center.
- B. Configure the Amazon CloudWatch agent on all servers within the local environment and publish metrics to Amazon CloudWatch Logs.
- C. Use AWS Application Discovery Service and enable agentless discovery in the existing virtualization environment.
- D. Enable AWS Application Discovery Service in the AWS Management Console and configure the corporate firewall to allow scans over a VPN.

**Answer:** A

**NEW QUESTION 109**

A Solutions Architect must create a cost-effective backup solution for a company's 500MB source code repository of proprietary and sensitive applications. The repository runs on Linux and backs up daily to tape. Tape backups are stored for 1 year.

The current solutions are not meeting the company's needs because it is a manual process that is prone to error, expensive to maintain, and does not meet the need for a Recovery Point Objective (RPO) of 1 hour or Recovery Time Objective (RTO) of 2 hours. The new disaster recovery requirement is for backups to be stored offsite and to be able to restore a single file if needed.

Which solution meets the customer's needs for RTO, RPO, and disaster recovery with the LEAST effort and expense?

- A. Replace local tapes with an AWS Storage Gateway virtual tape library to integrate with current backup software.
- B. Run backups nightly and store the virtual tapes on Amazon S3 standard storage in US-EAST-1. Use cross-region replication to create a second copy in US-WEST-2. Use Amazon S3 lifecycle policies to perform automatic migration to Amazon Glacier and deletion of expired backups after 1 year?
- C. Configure the local source code repository to synchronize files to an AWS Storage Gateway file Amazon gateway to store backup copies in an Amazon S3 Standard bucket.

- D. Enable versioning on the Amazon S3 bucket
- E. Create Amazon S3 lifecycle policies to automatically migrate old versions of objects to Amazon S3 Standard 0 Infrequent Access, then Amazon Glacier, then delete backups after 1 year.
- F. Replace the local source code repository storage with a Storage Gateway stored volume
- G. Change the default snapshot frequency to 1 hour
- H. Use Amazon S3 lifecycle policies to archive snapshots to Amazon Glacier and remove old snapshots after 1 year
- I. Use cross-region replication to create a copy of the snapshots in US-WEST-2.
- J. Replace the local source code repository storage with a Storage Gateway cached volume
- K. Create a snapshot schedule to take hourly snapshots
- L. Use an Amazon CloudWatch Events schedule expression rule to run on hourly AWS Lambda task to copy snapshots from US-EAST -1 to US-WEST-2.

**Answer: B**

**Explanation:**

<https://d1.awsstatic.com/whitepapers/aws-storage-gateway-file-gateway-for-hybrid-architectures.pdf>

**NEW QUESTION 111**

A company has more than 100 AWS accounts, with one VPC per account, that need outbound HTTPS connectivity to the internet. The current design contains one NAT gateway per Availability Zone (AZ) in each VPC. To reduce costs and obtain information about outbound traffic, management has asked for a new architecture for internet access.

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

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

**Answer: D**

**Explanation:**

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

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

**NEW QUESTION 112**

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 keys
- 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 keys
- 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 keys
- G. Use the CloudHSM 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 keys
- 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 114**

A company operating a website on AWS requires high levels of scalability, availability and performance. The company is running a Ruby on Rails application on Amazon EC2. It has a data tier on MySQL 5.6 on Amazon EC2 using 16 TB of Amazon EBS storage. Amazon CloudFront is used to cache application content. The Operations team is reporting continuous and unexpected growth of EBS volumes assigned to the MySQL database. The Solutions Architect has been asked to design a highly scalable, highly available, and high-performing solution.

Which solution is the MOST cost-effective at scale?

- A. Implement Multi-AZ and Auto Scaling for all EC2 instances in the current configuration
- B. Ensure that all EC2 instances are purchased as reserved instances
- C. Implement new elastic Amazon EBS volumes for the data tier.
- D. Design and implement the Docker-based containerized solution for the application using Amazon EC2
- E. Migrate to an Amazon Aurora MySQL Multi-AZ cluster
- F. Implement storage checks for Aurora MySQL storage utilization and an AWS Lambda function to grow the Aurora MySQL storage, as necessary
- G. Ensure that Multi-AZ architectures are implemented.
- H. Ensure that EC2 instances are right-sized and behind an Elastic Load Balancing load balancer. Implement Auto Scaling with EC2 instances
- I. Ensure that the reserved instances are purchased for fixed capacity and that Auto Scaling instances run on demand
- J. Migrate to an Amazon Aurora MySQL Multi-AZ cluster
- K. Ensure that Multi-AZ architectures are implemented.
- L. Ensure that EC2 instances are right-sized and behind an Elastic Load Balance

- M. Implement Auto Scaling with EC2 instance
- N. Ensure that Reserved instances are purchased for fixed capacity and that Auto Scaling instances run on demand
- O. Migrate to an Amazon Aurora MySQL Multi-AZ cluster
- P. Implement storage checks for Aurora MySQL storage utilization and an AWS Lambda function to grow Aurora MySQL storage, as necessary
- Q. Ensure Multi-AZ architectures are implemented.

**Answer: C**

#### NEW QUESTION 115

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 118

A company has an internal AWS Elastic Beanstalk worker environment inside a VPC that must access an external payment gateway API available on an HTTPS endpoint on the public internet. Because of security policies, the payment gateway's Application team can grant access to only one public IP address.

Which architecture will set up an Elastic Beanstalk environment to access the company's application without making multiple changes on the company's end?

- A. Configure the Elastic Beanstalk application to place Amazon EC2 instances in a private subnet with an outbound route to a NAT gateway in a public subnet. Associate an Elastic IP address to the NAT gateway that can be whitelisted on the payment gateway application side
- B. Configure the Elastic Beanstalk application to place Amazon EC2 instances in a public subnet with an internet gateway. Associate an Elastic IP address to the internet gateway that can be whitelisted on the payment gateway application side
- C. Configure the Elastic Beanstalk application to place Amazon EC2 instances in a private subnet. Set an `https_proxy` application parameter to send outbound HTTPS connections to an EC2 proxy server deployed in a public subnet. Associate an Elastic IP address to the EC2 proxy host that can be whitelisted on the payment gateway application side
- D. Configure the Elastic Beanstalk application to place Amazon EC2 instances in a public subnet. Set the `https_proxy` and `no_proxy` application parameters to send non-VPC outbound HTTPS connections to an EC2 proxy server deployed in a public subnet. Associate an Elastic IP address to the EC2 proxy host that can be whitelisted on the payment gateway application side

**Answer: C**

#### NEW QUESTION 123

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 126**

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

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

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

**Answer:** BDE

**NEW QUESTION 127**

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/awsccloudtrail/latest/userguide/best-practices-security.html](https://docs.aws.amazon.com/en_pv/awsccloudtrail/latest/userguide/best-practices-security.html)

**NEW QUESTION 128**

An organization has a write-intensive mobile application that uses Amazon API Gateway, AWS Lambda, and Amazon DynamoDB. The application has scaled well, however, costs have increased exponentially because of higher than anticipated Lambda costs. The application's use is unpredictable, but there has been a steady 20% increase in utilization every month.

While monitoring the current Lambda functions, the Solutions Architect notices that the execution-time averages 4.5 minutes. Most of the wait time is the result of a high-latency network call to a 3-TB MySQL database server that is on-premises. A VPN is used to connect to the VPC, so the Lambda functions have been configured with a five-minute timeout.

How can the Solutions Architect reduce the cost of the current architecture?

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

**Answer:** D

**NEW QUESTION 129**

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

- \* The application code will need a 500 MB static dataset to be present before application startup.
- \* The application tier be able to scale Up and down based on demand with as little startup time as possible.
- \* The development team should be able to update the code multiple times each day.
- \* Critical operating system (OS) patches must be installed within 48 hours of being released. Which deployment strategy meets these requirements?

- A. Use AWS Manager to create a new AMI with the updated OS patches. Update the Auto Scaling group to use the patches AMI and replace existing unpatched

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

**Answer: C**

#### NEW QUESTION 130

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

Key requirements are:

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

A key requirement is that the environment has no access to the internet, either directly or via the on-premises proxy. However, the application needs to be able to seamlessly communicate to Amazon S3, Amazon DynamoDB, and Amazon EC2 API, without the need for reconfiguration for each new deployment.

Which of the following should the Solutions Architect do to achieve this target architecture? (Choose three.)

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

**Answer: ACE**

#### Explanation:

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

#### NEW QUESTION 135

To abide by industry regulations, a Solutions Architect must design a solution that will store a company's critical data in multiple public AWS Regions, including in the United States, where the company's headquarters is located. The Solutions Architect is required to provide access to the data stored in AWS to the company's global WAN network. The Security team mandates that no traffic accessing this data should traverse the public internet.

How should the Solutions Architect design a highly available solution that meets the requirements and is cost-effective?

- A. Establish AWS Direct Connect connections from the company headquarters to all AWS Regions in use. Use the company WAN to send traffic over to the headquarters and then to the respective DX connection to access the data.
- B. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connectio
- C. Use inter-region VPC peering to access the data in other AWS Regions.
- D. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connectio
- E. Use an AWS transit VPC solution to access data in other AWS Regions.
- F. Establish two AWS Direct Connect connections from the company headquarters to an AWS Region. Use the company WAN to send traffic over a DX connectio
- G. Use Direct Connect Gateway to access data in other AWS Regions.

**Answer: D**

#### Explanation:

<https://aws.amazon.com/blogs/aws/new-aws-direct-connect-gateway-inter-region-vpc-access/>

#### NEW QUESTION 136

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

Which option meets these requirements?

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

**Answer: A**

#### Explanation:

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

**NEW QUESTION 137**

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

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

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

- A. Create accounts for each production workload within an organization in AWS Organizations. Place the production accounts within an organizational unit (OU). For each account, delete the default VPC. Create an SCP with a Deny rule for the attach an internet gateway and create a default VPC. Attach the SCP to the OU for the production accounts.
- B. Create accounts for each production workload within an organization in AWS Organizations. Place the production accounts within an organizational unit (OU). Create an SCP with a Deny rule on the attach an internet gateway action. Create an SCP with a Deny rule to prevent use of the default VPC. Attach the SCPs to the OU for the production accounts.
- C. Create a SCP containing a Deny Effect for cloudfront. 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 139**

A company is running a web application with On-Demand Amazon EC2 instances in Auto Scaling groups that scale dynamically based on custom metrics. After extensive testing, the company determines that the m5.2xlarge instance size is optimal for the workload. Application data is stored in db.r4.xlarge Amazon RDS instances that are confirmed to be optimal. The traffic to the web application spikes randomly during the day.

What other cost-optimization methods should the company implement to further reduce costs without impacting the reliability of the application?

- A. Double the instance count in the Auto Scaling groups and reduce the instance size to m5.large
- B. Reserve capacity for the RDS database and the minimum number of EC2 instances that are constantly running
- C. Reduce the RDS instance size to db.r4.xlarge and add five equivalent-sized read replicas to provide reliability
- D. Reserve capacity for all EC2 instances and leverage Spot Instance pricing for the RDS database

**Answer:** B

**NEW QUESTION 143**

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 147**

A company is using an Amazon CloudFront distribution to distribute both static and dynamic content from a web application running behind an Application Load Balancer. The web application requires user authorization and session tracking for dynamic content. The CloudFront distribution has a single cache behavior configured to forward the Authorization, Host, and User-Agent HTTP whitelist headers and a session cookie to the origin. All other cache behavior settings are set to their default value.

A valid ACM certificate is applied to the CloudFront distribution with a matching CNAME in the distribution settings. The ACM certificate is also applied to the HTTPS listener for the Application Load Balancer. The CloudFront origin protocol policy is set to HTTPS only. Analysis of the cache statistics report shows that the miss rate for this distribution is very high.

What can the Solutions Architect do to improve the cache hit rate for this distribution without causing the SSL/TLS handshake between CloudFront and the Application Load Balancer to fail?

- A. Create two cache behaviors for static and dynamic content
- B. Remove the User-Agent and Host HTTP headers from the whitelist headers section on both of the cache behaviors
- C. Remove the session cookie from the whitelist cookies section and the Authorization HTTP header from the whitelist headers section for the cache behavior configured for static content.
- D. Remove the User-Agent and Authorization HTTP headers from the whitelist headers section of the cache behavior
- E. Then update the cache behavior to use presigned cookies for authorization.
- F. Remove the Host HTTP header from the whitelist headers section and remove the session cookie from the whitelist cookies section for the default cache

behavior

G. Enable automatic object compression and use Lambda@Edge viewer request events for user authorization.

H. Create two cache behaviors for static and dynamic content

I. Remove the User-Agent HTTP header from the whitelist headers section on both of the cache behavior

J. Remove the session cookie from the whitelist cookies section and the Authorization HTTP header from the whitelist headers section for cache behavior configured for static content.

**Answer:** D

#### NEW QUESTION 148

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

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

A. Amazon Rekognition to identify who is calling.

B. Amazon Connect to create a cloud-based contact center.

C. Amazon Alexa for Business to build conversational interface.

D. AWS Lambda to integrate with internal systems.

E. Amazon Lex to recognize the intent of the caller.

F. Amazon SQS to add incoming callers to a queue.

**Answer:** BDE

#### NEW QUESTION 151

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 153

A large company is migrating its entire IT portfolio to AWS. Each business unit in the company has a standalone AWS account that supports both development and test environments. New accounts to support production workloads will be needed soon.

The Finance department requires a centralized method for payment but must maintain visibility into each group's spending to allocate costs.

The Security team requires a centralized mechanism to control IAM usage in all the company's accounts. What combination of the following options meet the company's needs with LEAST effort? (Choose two.)

A. Use a collection of parameterized AWS CloudFormation templates defining common IAM permissions that are launched into each account

B. Require all new and existing accounts to launch the appropriate stacks to enforce the least privilege model.

C. Use AWS Organizations to create a new organization from a chosen payer account and define an organizational unit hierarchy

D. Invite the existing accounts to join the organization and create new accounts using Organizations.

E. Require each business unit to use its own AWS account

F. Tag each AWS account appropriately and enable Cost Explorer to administer chargebacks.

G. Enable all features of AWS Organizations and establish appropriate service control policies that filter IAM permissions for sub-accounts.

H. Consolidate all of the company's AWS accounts into a single AWS account

I. Use tags for billing purposes and IAM's Access Advice feature to enforce the least privilege model.

**Answer:** BD

#### NEW QUESTION 154

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 155

A company has deployed an application to multiple environments in AWS, including production and testing. The company has separate accounts for production and testing, and users are allowed to create additional application users for team members or services, as needed. The Security team has asked the Operations team for better isolation between production and testing with centralized controls on security credentials and improved management of permissions between environments.

Which of the following options would MOST securely accomplish this goal?

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

**Answer:** A

**Explanation:**

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

#### NEW QUESTION 156

A large global company wants to migrate a stateless mission-critical application to AWS. The application is based on IBM WebSphere (application and integration middleware), IBM MQ (messaging middleware), and IBM DB2 (database software) on a z/OS operating system.

How should the Solutions Architect migrate the application to AWS?

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

**Answer:** B

**Explanation:**

<https://aws.amazon.com/blogs/database/aws-database-migration-service-and-aws-schema-conversion-tool-now->

<https://aws.amazon.com/quickstart/architecture/ibm-mq/>

#### NEW QUESTION 160

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

Which design meets the required request rate and response time?

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

**Answer:** C

**Explanation:**

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

#### NEW QUESTION 164

An auction website enables users to bid on collectible items. The auction rules require that each bid is processed only once and in the order it was received. The current implementation is based on a fleet of Amazon EC2 web servers that write bid records into Amazon Kinesis Data Streams. A single t2.large instance has a cron job that runs the bid processor, which reads incoming bids from Kinesis Data Streams and processes each bid. The auction site is growing in popularity, but users are complaining that some bids are not registering.

Troubleshooting indicates that the bid processor is too slow during peak demand hours, sometimes crashes while processing, and occasionally loses track of which records are being processed.

What changes should make the bid processing more reliable?

- A. Refactor the web application to use the Amazon Kinesis Producer Library (KPL) when posting bids to Kinesis Data Stream
- B. Refactor the bid processor to flag each record in Kinesis Data Streams as being unread, processing, and processed
- C. At the start of each bid processing run, scan Kinesis Data Streams for unprocessed records.
- D. Refactor the web application to post each incoming bid to an Amazon SNS topic in place of Kinesis Data Stream
- E. Configure the SNS topic to trigger an AWS Lambda function that processes each bid as soon as a user submits it.
- F. Refactor the web application to post each incoming bid to an Amazon SQS FIFO queue in place of Kinesis Data Stream
- G. Refactor the bid processor to continuously poll the SQS queue
- H. Place the bid processing EC2 instance in an Auto Scaling group with a minimum and a maximum size of 1.
- I. Switch the EC2 instance type from t2.large to a larger general compute instance type
- J. Put the bid processor EC2 instances in an Auto Scaling group that scales out the number of EC2 instances running the bid processor, based on the IncomingRecords metric in Kinesis Data Streams.

**Answer: C**

**Explanation:**

FIFO is better in this case compared to Kinesis, as it guarantees the order of the bid. Min Max 1, is okay as the SQS will hold the queue in case of failure of the instance, till it comes back again.

**NEW QUESTION 168**

A company with multiple accounts is currently using a configuration that does not meet the following security governance policies

- Prevent ingress from port 22 to any Amazon EC2 instance
- Require billing and application tags for resources
- Encrypt all Amazon EBS volumes

A Solutions Architect wants to provide preventive and detective controls including notifications about a specific resource, if there are policy deviations. Which solution should the Solutions Architect implement?

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

**Answer: C**

**NEW QUESTION 170**

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](https://docs.aws.amazon.com/IAM/latest/UserGuide/access_iam-tags.html)

**NEW QUESTION 174**

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 177**

.....

## 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>