Exam Questions SAP-C01

AWS Certified Solutions Architect- Professional

https://www.2passeasy.com/dumps/SAP-C01/

NEW QUESTION 1
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 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 2
An organization has two Amazon EC2 instances:

> The first is running an ordering application and an inventory application.

> The second is running a queuing system.
During certain times of the year, several thousand orders are placed per second. Some orders were lost when the queuing system was down. Also, the organization's inventory application has the incorrect quantity of products because some orders were processed twice.
What should be done to ensure that the applications can handle the increasing number of orders?

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

Answer: C

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

NEW QUESTION 3
A company runs a memory-intensive analytics application using on-demand Amazon EC2 compute optimized instance. The application is used continuously and application demand doubles during working hours. The application currently scales based on CPU usage. When scaling in occurs, a lifecycle hook is used because the instance requires 4 minutes to clean the application state before terminating.
Because users reported poor performance during working hours, scheduled scaling actions were implemented so additional instances would be added during working hours. The Solutions Architect has been asked to reduce the cost of the application.
Which solution is MOST cost-effective?

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

Answer: D

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

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

A. Create an Elasticsearch server on Amazon EC2 right-sized with 2 TB of Amazon EBS and a public AWS Elastic Beanstalk environment for the web applicatio

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

Answer: D

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

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

Answer: B

NEW QUESTION 6
A company has a requirement that only allows specially hardened AMIs to be launched into public subnets in a VPC, and for the AMIs to be associated with a specific security group. Allowing non-compliant instances to launch into the public subnet could present a significant security risk if they are allowed to operate.
A mapping of approved AMIs to subnets to security groups exists in an Amazon DynamoDB table in the same AWS account. The company created an AWS Lambda function that, when invoked, will terminate a given Amazon EC2 instance if the combination of AMI, subnet, and security group are not approved in the DynamoDB table.
What should the Solutions Architect do to MOST quickly mitigate the risk of compliance deviations?

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

Answer: C

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

NEW QUESTION 7
As a part of building large applications in the AWS Cloud, the Solutions Architect is required to implement the perimeter security protection. Applications running on AWS have the following endpoints:

> Application Load Balancer

> Amazon API Gateway regional endpoint

> Elastic IP address-based EC2 instances.

> Amazon S3 hosted websites.

> Classic Load Balancer

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

> DDoS protection

> SQL injection protection

> IP address whitelist/blacklist

> HTTP flood protection

> Bad bot scraper protection

How should the Solutions Architect design the solution?

A. Deploy AWS WAF and AWS Shield Advanced on all web endpoint
B. Add AWS WAF rules to enforce the company's requirements.

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

Answer: C

NEW QUESTION 8
A company is moving a business-critical, multi-tier application to AWS. The architecture consists of a desktop client application and server infrastructure. The server infrastructure resides in an on-premises data center that frequently fails to maintain the application uptime SLA of 99.95%. A Solutions Architect must re-architect the application to ensure that it can meet or exceed the SLA.
The application contains a PostgreSQL database running on a single virtual machine. The business logic and presentation layers are load balanced between multiple virtual machines. Remote users complain about slow load times while using this latency-sensitive application.
Which of the following will meet the availability requirements with little change to the application while improving user experience and minimizing costs?

A. Migrate the database to a PostgreSQL database in Amazon EC2. Host the application and presentation layers in automatically scaled Amazon ECS containers behind an Application Load Balance
B. Allocate an Amazon WorkSpaces WorkSpace for each end user to improve the user experience.
C. Migrate the database to an Amazon RDS Aurora PostgreSQL configuratio
D. Host the application and presentation layers in an Auto Scaling configuration on Amazon EC2 instances behind an Application Load Balance
E. Use Amazon AppStream 2.0 to improve the user experience.
F. Migrate the database to an Amazon RDS PostgreSQL Multi-AZ configuratio
G. Host the application andpresentation layers in automatically scaled AWS Fargate containers behind a Network Load Balance
H. Use Amazon ElastiCache to improve the user experience.
I. Migrate the database to an Amazon Redshift cluster with at least two node
J. Combine and host the application and presentation layers in automatically scaled Amazon ECS containers behind an Application Load Balance
K. Use Amazon CloudFront to improve the user experience.

Answer: B

NEW QUESTION 9
A company has an existing on-premises three-tier web application. The Linux web servers serve content from a centralized file share on a NAS server because the content is refreshed several times a day from various sources. The existing infrastructure is not optimized and the company would like to move to AWS in order to gain the ability to scale resources up and down in response to load. On-premises and AWS resources are connected using AWS Direct Connect.
How can the company migrate the web infrastructure to AWS without delaying the content refresh process?

A. Create a cluster of web server Amazon EC2 instances behind a Classic Load Balancer on AW
B. Share an Amazon EBS volume among all instances for the conten
C. Schedule a periodic synchronization of this volume and the NAS server.
D. Create an on-premises file gateway using AWS Storage Gateway to replace the NAS server and replicate content to AW
E. On the AWS side, mount the same Storage Gateway bucket to each web server Amazon EC2 instance to serve the content.
F. Expose an Amazon EFS share to on-premises users to serve as the NAS serv
G. Mount the same EFS share to the web server Amazon EC2 instances to serve the content.
H. Create web server Amazon EC2 instances on AWS in an Auto Scaling grou
I. Configure a nightly process where the web server instances are updated from the NAS server.

Answer: C

Explanation:
File gateway is limited by performance its gateway instance, whether EC2 or On-premises, Cache will get filled up fast if not properly configured, For large number of EC2 instances EFS scales better. So, bottom line is File Storage gateway is for legacy applications and you have to add cost of large gateway instances before comparing it to same quantity of EFS storage. https://www.reddit.com/r/aws/comments/82pyop/storage_gateway_vs_efs/
https://docs.aws.amazon.com/efs/latest/ug/efs-onpremises.html

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 rebuild other applications native to the cloud. The team created both Amazon EC2 databases and used Amazon RDS. Each team in the company was responsible for migrating their applications, and they have created individual accounts for isolation of resources. The company did not have much time to consider costs, but now it would like suggestions on reducing its AWS spend.
Which steps should a Solutions Architect take to reduce costs?

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

N. Have each team upload their bill to an Amazon S3 bucket for analysis of team spendin
O. Use Spot instances on nightly batch processing jobs.

Answer: B

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

NEW QUESTION 10
A three-tier web application runs on Amazon EC2 instances. Cron daemons are used to trigger scripts that collect the web server, application, and database logs and send them to a centralized location every hour. Occasionally, scaling events or unplanned outages have caused the instances to stop before the latest logs were collected, and the log files were lost.
Which of the following options is the MOST reliable way of collecting and preserving the log files?

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

Answer: C

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

NEW QUESTION 15
A company has multiple AWS accounts hosting IT applications. An Amazon CloudWatch Logs agent is installed on all Amazon EC2 instances. The company wants to aggregate all security events in a centralized AWS account dedicated to log storage.
Security Administrators need to perform near-real-time gathering and correlating of events across multiple AWS accounts.
Which solution satisfies these requirements?

A. Create a Log Audit IAM role in each application AWS account with permissions to view CloudWatch Logs, configure an AWS Lambda function to assume the Log Audit role, and perform an hourly export of CloudWatch Logs data to an Amazon S3 bucket in the logging AWS account.
B. Configure CloudWatch Logs streams in each application AWS account to forward events to CloudWatch Logs in the logging AWS accoun
C. In the logging AWS account, subscribe an Amazon Kinesis Data Firehose stream to Amazon CloudWatch Events, and use the stream to persist log data in Amazon S3.
D. Create Amazon Kinesis Data Streams in the logging account, subscribe the stream to CloudWatch Logs streams in each application AWS account, configure an Amazon Kinesis Data Firehose delivery stream with the Data Streams as its source, and persist the log data in an Amazon S3 bucket inside the logging AWS account.
E. Configure CloudWatch Logs agents to publish data to an Amazon Kinesis Data Firehose stream in the logging AWS account, use an AWS Lambda function to read messages from the stream and push messages to Data Firehose, and persist the data in Amazon S3.

Answer: C

Explanation:
The solution uses Amazon Kinesis Data Streams and a log destination to set up an endpoint in the logging account to receive streamed logs and uses Amazon Kinesis Data Firehose to deliver log data to the Amazon Simple Storage Solution (S3) bucket. Application accounts will subscribe to stream all (or part) of their Amazon CloudWatch logs to a defined destination in the logging account via subscription filters. https://aws.amazon.com/blogs/architecture/central-logging-in-multi-account-environments/

NEW QUESTION 18
A company is migrating an application to AWS. It wants to use fully managed services as much as possible during the migration. The company needs to store large, important documents within the application with the following requirements:

≫ The data must be highly durable and available.

≫ The data must always be encrypted at rest and in transit.

≫ The encryption key must be managed by the company and rotated periodically. Which of the following solutions should the Solutions Architect recommend?

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

Answer: B

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

NEW QUESTION 21
A company runs a dynamic mission-critical web application that has an SLA of 99.99%. Global application users access the application 24/7. The application is currently hosted on premises and routinely fails to meet its SLA, especially when millions of users access the application concurrently. Remote users complain of

latency.
How should this application be redesigned to be scalable and allow for automatic failover at the lowest cost?

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

Answer: C

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

NEW QUESTION 23
A retail company has a custom NET web application running on AWS that uses Microsoft SQL Server for the database The application servers maintain a user's session locally.
Which combination of architecture changes are needed ensure all tiers of the solution are highly available? (Select THREE.)

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

Answer: BDE

NEW QUESTION 28
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 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 regio
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 regio
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 mod
F. Place the web and the application tiers in an Auto Scaling behind a load balancer, which can automatically scale when the load arrives to the applicatio
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 capacit
I. Activate the primary database in one region only and the standby database in the other regio
J. Use Amazon Route 53 to automatically switch traffic from one region to another using health check routing policies.

Answer: C

NEW QUESTION 32
A Solutions Architect has created an AWS CloudFormation template for a three-tier application that contains an Auto Scaling group of Amazon EC2 instances running a custom AMI.
The Solutions Architect wants to ensure that future updates to the custom AMI can be deployed to a running stack by first updating the template to refer to the new AMI, and then invoking UpdateStack to replace the EC2 instances with instances launched from the new AMI.
How can updates to the AMI be deployed to meet these requirements?

A. Create a change set for a new version of the template, view the changes to the running EC2 instances to ensure that the AMI is correctly updated, and then execute the change set.
B. Edit the AWS::AutoScaling::LaunchConfiguration resource in the template, changing its DeletionPolicy to Replace.
C. Edit the AWS::AutoScaling:: AutoScalingGroup resource in the template, inserting an UpdatePolicy attribute.
D. Create a new stack from the updated templat
E. Once it is successfully deployed, modify the DNS records to point to the new stack and delete the old stack.

Answer: C

Explanation:

References:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-properties-as-launchconfig.html

NEW QUESTION 35
A company is adding a new approved external vendor that only supports IPv6 connectivity. The company's backend systems sit in the private subnet of an Amazon VPC. The company uses a NAT gateway to allow these systems to communicate with external vendors over IPv4. Company policy requires systems that communicate with external vendors use a security group that limits access to only approved external vendors. The virtual private cloud (VPC) uses the default network ACL.
The Systems Operator successfully assigns IPv6 addresses to each of the backend systems. The Systems Operator also updates the outbound security group to include the IPv6 CIDR of the external vendor (destination). The systems within the VPC are able to ping one another successfully over IPv6. However, these systems are unable to communicate with the external vendor.
What changes are required to enable communication with the external vendor?

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

Answer: D

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

NEW QUESTION 38
A company is using AWS to run an internet-facing production application written in Node.js. The Development team is responsible for pushing new versions of their software directly to production. The application software is updated multiple times a day. The team needs guidance from a Solutions Architect to help them deploy the software to the production fleet quickly and with the least amount of disruption to the service.
Which option meets these requirements?

A. Prepackage the software into an AMI and then use Auto Scaling to deploy the production flee
B. For software changes, update the AMI and allow Auto Scaling to automatically push the new AMI to production.
C. Use AWS CodeDeploy to push the prepackaged AMI to productio
D. For software changes, reconfigure CodeDeploy with new AMI identification to push the new AMI to the production fleet.
E. Use AWS Elastic Beanstalk to host the production applicatio
F. For software changes, upload the new application version to Elastic Beanstalk to push this to the production fleet using a blue/green deployment method.
G. Deploy the base AMI through Auto Scaling and bootstrap the software using user dat
H. For software changes, SSH to each of the instances and replace the software with the new version.

Answer: C

NEW QUESTION 41
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 distributio
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 44
AnyCompany has acquired numerous companies over the past few years. The CIO for AnyCompany would like to keep the resources for each acquired company separate. The CIO also would like to enforce a chargeback model where each company pays for the AWS services it uses.
The Solutions Architect is tasked with designing an AWS architecture that allows AnyCompany to achieve the following:

> Implementing a detailed chargeback mechanism to ensure that each company pays for the resources it uses.

> AnyCompany can pay for AWS services for all its companies through a single invoice.

> Developers in each acquired company have access to resources in their company only.

> Developers in an acquired company should not be able to affect resources in their company only.

> A single identity store is used to authenticate Developers across all companies.
Which of the following approaches would meet these requirements? (Choose two.)

A. Create a multi-account strategy with an account per compan
B. Use consolidated billing to ensure that AnyCompany needs to pay a single bill only.
C. Create a multi-account strategy with a virtual private cloud (VPC) for each compan
D. Reduce impact across companies by not creating any VPC peering link

E. As everything is in a single account, there will be a single invoic

F. use tagging to create a detailed bill for each company.

G. Create IAM users for each Developer in the account to which they require acces

H. Create policies that allow the users access to all resources in that accoun

I. Attach the policies to the IAM user.

J. Create a federated identity store against the company's Active Director

K. Create IAM roles with appropriate permissions and set the trust relationships with AWS and the identity stor

L. Use AWS STS to grant users access based on the groups they belong to in the identity store.

M. Create a multi-account strategy with an account per compan

N. For billing purposes, use a tagging solution that uses a tag to identify the company that creates each resource.

Answer: AD

NEW QUESTION 47
A Solutions Architect must establish a patching plan for a large mixed fleet of Windows and Linux servers. The patching plan must be implemented securely, be audit ready, and comply with the company's business requirements.
Which option will meet these requirements with MINIMAL effort?

A. Install and use an OS-native patching service to manage the update frequency and release approval for all instance

B. Use AWS Config to verify the OS state on each instance and report on any patch compliance issues.

C. Use AWS Systems Manager on all instances to manage patchin

D. Test patches outside of production and then deploy during a maintenance window with the appropriate approval.

E. Use AWS OpsWorks for Chef Automate to run a set of scripts that will iterate through all instances of a given typ

F. Issue the appropriate OS command to get and install updates on each instance, including any required restarts during the maintenance window.

G. Migrate all applications to AWS OpsWorks and use OpsWorks automatic patching support to keep the OS up-to-date following the initial installatio

H. Use AWS Config to provide audit and compliance reporting.

Answer: B

Explanation:
Only Systems Manager can patch both OS effectively on AWS and on premise.

NEW QUESTION 49
A Solutions Architect must migrate an existing on-premises web application with 70 TB of static files supporting a public open-data initiative. The architect wants to upgrade to the latest version of the host operating system as part of the migration effort.
Which is the FASTEST and MOST cost-effective way to perform the migration?

A. Run a physical-to-virtual conversion on the application serve

B. Transfer the server image over the internet, and transfer the static data to Amazon S3.

C. Run a physical-to-virtual conversion on the application serve

D. Transfer the server image over AWS Direct Connect, and transfer the static data to Amazon S3.

E. Re-platform the server to Amazon EC2, and use AWS Snowball to transfer the static data to Amazon S3.

F. Re-platform the server by using the AWS Server Migration Service to move the code and data to a new Amazon EC2 instance.

Answer: C

NEW QUESTION 50
A company's main intranet page has experienced degraded response times as its user base has increased although there are no reports of users seeing error pages. The application uses Amazon DynamoDB in read-only mode.
Amazon DynamoDB latency metrics for successful requests have been in a steady state even during times when users have reported degradation The Development team has correlated the issue to ProvisionedThrough put Exceeded exceptions in the application logs when doing Scan and read operations The team also identified an access pattern of steady spikes of read activity on a distributed set of individual data items
The Chief Technology Officer wants to improve the user experience
Which solutions will meet these requirements with the LEAST amount of changes to the application? (Select TWO)

A. Change the data model of the DynamoDB tables to ensure that all Scan and read operations meet DynamoDB best practices of uniform data access, reaching the full request throughput provisioned for the DynamoDB tables

B. Enable DynamoDB auto scaling to manage the throughput capacity as table traffic increases Set the upper and lower limits to control costs and set a target utilization given the peak usage and how quickly the traffic changes.

C. Provision Amazon ElastiCache for Redis with cluster mode enabled The cluster should be provisioned with enough shards to spread the application load and provision at least one read replica node for each shard

D. Implement the DynamoDB Accelerator (DAX) client and provision a DAX cluster with the appropriate node types to sustain the application loa

E. Tune the item and query cache configuration for an optimal user experience

F. Remove error retries and exponential backoffs in the application code to handle throttling errors

Answer: AE

NEW QUESTION 51
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 configuratio

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 performe

F. Archive log data by using a batch exportto 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/awscloudtrail/latest/userguide/cloudwatch-alarms-for-cloudtrail.html
https://docs.aws.amazon.com/en_pv/awscloudtrail/latest/userguide/best-practices-security.html

NEW QUESTION 52
The CISO of a large enterprise with multiple IT departments, each with its own AWS account, wants one central place where AWS permissions for users can be managed and users authentication credentials can be synchronized with the company's existing on-premises solution.
Which solution will meet the CISO's requirements?

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

Answer: A

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

NEW QUESTION 53
A Solution Architect is designing a deployment strategy for an application tier and gas the following requirements.
* The application code will need a 500 HB static dataset to be present before application startup.
* The application tier be able to scale Up and down based on demand with as little startup time as possible.
* The development team should be able to update the code multiple times each day.
* Critical operating system (OS) patches must be installed within 48 hours of being released. Which deployment strategy meets these requirements?

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

Answer: C

NEW QUESTION 55
A company is currently using AWS CodeCommit for its source control and AWS CodePipeline for continuous integration. The pipeline has a build stage for building the artifacts which is then staged in an Amazon S3 bucket.
The company has identified various improvement opportunities in the existing process, and a Solutions Architect has been given the following requirement:

> Create a new pipeline to support feature development

> Support feature development without impacting production applications

> Incorporate continuous testing with unit tests

> Isolate development and production artifacts

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

A. Trigger a separate pipeline from CodeCommit feature branche
B. Use AWS CodeBuild for running unit test
C. Use CodeBuild to stage the artifacts within an S3 bucket in a separate testing account.
D. Trigger a separate pipeline from CodeCommit feature branche
E. Use AWS Lambda for running unit test
F. Use AWS CodeDeploy to stage the artifacts within an S3 bucket in a separate testing account.
G. Trigger a separate pipeline from CodeCommit tags Use Jenkins for running unit test

H. Create a stage in the pipeline with S3 as the target for staging the artifacts with an S3 bucket in a separate testing account.
I. Create a separate CodeCommit repository for feature development and use it to trigger the pipelin
J. Use AWS Lambda for running unit test
K. Use AWS CodeBuild to stage the artifacts within different S3 buckets in the same production account.

Answer: A

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

NEW QUESTION 58
A utility company wants to collect usage data every 5 minutes from its smart meters to facilitate time-of-use metering When a meter sends data to AWS the data is sent to Amazon API Gateway, processed by an AWS Lambda function and stored in an Amazon DynamoDB table During the pilot phase, the Lambda functions took from 3 to 5 seconds to complete
As more smart meters are deployed, the Engineers notice the Lambda functions are taking from 1 to 2 minutes to complete The functions are also increasing in duration as new types of metrics are collected from the devices There are many ProvisionedThroughputExceededException errors while performing PUT operations on DynamoDB and there are also many TooMany Requests Exception errors from Lambda.
Which combination of changes will resolve these issues? (Select TWO)

A. increase the write capacity units to the DynamoDB table
B. Increase the memory available to the Lambda functions
C. Increase the payload size from the smart meters to send more data
D. Stream the data into an Amazon Kinesis data stream from API Gateway and process the data in batches
E. Collect data in an Amazon SQS FIFO queue, which triggers a Lambda function to process each message

Answer: AB

NEW QUESTION 59
A company is running a high-user-volume media-sharing application on premises It currently hosts about 400 TB of data with millions of video files The company is migrating this application to AWS to improve reliability and reduce costs
The Solutions Architecture team plans to store the videos in an Amazon S3 bucket and use Amazon
CloudFront to distribute videos to users. The company needs to migrate this application to AWS within 10 days with the least amount of downtime possible. The company currently has 1 Gbps connectivity to the internet with 30 percent free capacity
Which of the following solutions would enable the company to migrate the workload to AWS and meet an of the requirements?

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

Answer: D

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

NEW QUESTION 61
A company has implemented AWS Organizations. It has recently set up a number of new accounts and wants to deny access to a specific set of AWS services in these new accounts.
How can this be controlled MOST efficiently?

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

Answer: B

NEW QUESTION 66
A company is using an Amazon CloudFront distribution to distribute both static and dynamic content from a web application running behind an Application Load Balancer. The web application requires user authorization and session tracking for dynamic content. The CloudFront distribution has a single cache behavior configured to forward the Authorization, Host, and User-Agent HTTP whitelist headers and a session cookie to the origin. All other cache behavior settings are set to their default value.
A valid ACM certificate is applied to the CloudFront distribution with a matching CNAME in the distribution settings. The ACM certificate is also applied to the HTTPS listener for the Application Load Balancer. The CloudFront origin protocol policy is set to HTTPS only. Analysis of the cache statistics report shows that the miss rate for this distribution is very high.
What can the Solutions Architect do to improve the cache hit rate for this distribution without causing the SSL/TLS handshake between CloudFront and the Application Load Balancer to fail?

A. Create two cache behaviors for static and dynamic conten
B. Remove the User-Agent and Host HTTP headers from the whitelist headers section on both if the cache behavior
C. Remove the session cookie from the whitelist cookies section and the Authorization HTTP header from the whitelist headers section for cache behavior configured for static content.

D. Remove the User-Agent and Authorization HTTP headers from the whitelist headers section of the cache behavio
E. Then update the cache behavior to use presigned cookies for authorization.
F. Remove the Host HTTP header from the whitelist headers section and remove the session cookie from the whitelist cookies section for the default cache behavio
G. Enable automatic object compression and use Lambda@Edge viewer request events for user authorization.
H. Create two cache behaviors for static and dynamic conten
I. Remove the User-Agent HTTP header from the whitelist headers section on both of the cache behavior
J. Remove the session cookie from the whitelist cookies section and the Authorization HTTP header from the whitelist headers section for cache behavior configured for static content.

Answer: D

NEW QUESTION 69
A company has asked a Solutions Architect to design a secure content management solution that can be accessed by API calls by external customer applications. The company requires that a customer administrator must be able to submit an API call and roll back changes to existing files sent to the content management solution, as needed.
What is the MOST secure deployment design that meets all solution requirements?

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

Answer: A

NEW QUESTION 73
A company wants to migrate its website from an on-premises data center onto AWS. At the same time, it wants to migrate the website to a containerized microservice-based architecture to improve the availability and cost efficiency. The company's security policy states that privileges and network permissions must be configured according to best practice, using least privilege.
A Solutions Architect must create a containerized architecture that meets the security requirements and has deployed the application to an Amazon ECS cluster.
What steps are required after the deployment to meet the requirements? (Choose two.)

A. Create tasks using the bridge network mode.
B. Create tasks using the awsvpc network mode.
C. Apply security groups to Amazon EC2 instances, and use IAM roles for EC2 instances to access other resources.
D. Apply security groups to the tasks, and pass IAM credentials into the container at launch time to access other resources.
E. Apply security groups to the tasks, and use IAM roles for tasks to access other resources.

Answer: BE

Explanation:
https://aws.amazon.com/about-aws/whats-new/2017/11/amazon-ecs-introduces-awsvpc-networking-mode-for-c
https://amazonaws-china.com/blogs/compute/introducing-cloud-native-networking-for-ecs-containers/
https://docs.aws.amazon.com/AmazonECS/latest/developerguide/task-iam-roles.html

NEW QUESTION 78
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 accoun
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 hierarch
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 accoun
I. Use tags for billing purposes and IAM's Access Advice feature to enforce the least privilege model.

Answer: BD

NEW QUESTION 82
A company is having issues with a newly deployed server less infrastructure that uses Amazon API Gateway, Amazon Lambda, and Amazon DynamoDB.
In a steady state, the application performs as expected However, during peak load, tens of thousands of simultaneous invocations are needed and user request fail multiple times before succeeding. The company has checked the logs for each component, focusing specifically on Amazon CloudWatch Logs for Lambda.
There are no error logged by the services or applications.

What might cause this problem?

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

Answer: A

NEW QUESTION 84
During a security audit of a Service team's application a Solutions Architect discovers that a username and password tor an Amazon RDS database and a set of AWSIAM user credentials can be viewed in the AWS Lambda function code. The Lambda function uses the username and password to run queries on the database and it uses the I AM credentials to call AWS services in a separate management account.
The Solutions Architect is concerned that the credentials could grant inappropriate access to anyone who can view the Lambda code The management account and the Service team's account are in separate AWS Organizations organizational units (OUs)
Which combination of changes should the Solutions Architect make to improve the solution's security? (Select TWO)

A. Configure Lambda to assume a tole in the management account with appropriate access to AWS
B. Configure Lambda to use the stored database credentials in AWS Secrets Manager and enable automatic rotation
C. Create a Lambda function to rotate the credentials every hour by deploying a new Lambda version with the updated credentials
D. Use an SCP on the management accounts OU to prevent IAM users from accessing resources m the Service team's account
E. Enable AWS Shield Advanced on the management account to shield sensitive resources from unauthorized IAM access

Answer: BD

NEW QUESTION 88
A company has an application that runs a web service on Amazon EC2 instances and stores .jpg images in Amazon S3. The web traffic has a predictable baseline, but often demand spikes unpredictably for short periods of time. The application is loosely coupled and stateless. The .jpg images stored in Amazon S3 are accessed frequently for the first 15 to 20 days, they are seldom accessed thereafter but always need to be immediately available. The CIO has asked to find ways to reduce costs.
Which of the following options will reduce costs? (Choose two.)

A. Purchase Reserved instances for baseline capacity requirements and use On-Demand instances for the demand spikes.
B. Configure a lifecycle policy to move the .jpg images on Amazon S3 to S3 IA after 30 days.
C. Use On-Demand instances for baseline capacity requirements and use Spot Fleet instances for the demand spikes.
D. Configure a lifecycle policy to move the .jpg images on Amazon S3 to Amazon Glacier after 30 days.
E. Create a script that checks the load on all web servers and terminates unnecessary On-Demand instances.

Answer: AB

NEW QUESTION 91
A company with multiple accounts is currently using a configuration that does not meet the following security governance policies
• Prevent ingress from port 22 to any Amazon EC2 instance
• Require billing and application tags for resources
• Encrypt all Amazon EBS volumes
A Solutions Architect wants to provide preventive and detective controls including notifications about a specific resource, if there are policy deviations.
Which solution should the Solutions Architect implement?

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

Answer: C

NEW QUESTION 96
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual SAP-C01 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the SAP-C01 Product From:

https://www.2passeasy.com/dumps/SAP-C01/

Money Back Guarantee

SAP-C01 Practice Exam Features:

* SAP-C01 Questions and Answers Updated Frequently

* SAP-C01 Practice Questions Verified by Expert Senior Certified Staff

* SAP-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* SAP-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year