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 syste
D. Mount this file system to all 10 instances.
E. Create a new Amazon S3 bucket and move all the data to this new bucke
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 servic
B. Set up an Amazon ECS service that is fronted by an Application Load Balancer (ALB). Use a custom authenticator to control access to the AP
C. Store request meta information in Amazon DynamoDB with Auto Scaling and static content in a secured S3 bucke
D. Make secure signed requests for Amazon S3 objects and proxy the data through the REST service interface.
E. Use AWS Fargate to host a container that runs a self-contained REST servic
F. Set up an ECS service that is fronted by a cross-zone AL
G. Use an Amazon Cognito user pool to control access to the AP
H. Store request meta information in DynamoDB with Auto Scaling and static content in a secured S3 bucke
I. Generate presigned URLs when returning references to content stored in Amazon S3.
J. Set up Amazon API Gateway and create the required API resources and method
K. Use an Amazon Cognito user pool to control access to the AP
L. Configure the methods to use AWS Lambda proxy integrations, and process each resource with a unique AWS Lambda functio
M. Store request meta information in DynamoDB with Auto Scaling and static content in a secured S3 bucke
N. Generate presigned URLs when returning references to content stored in Amazon S3.
O. Set up Amazon API Gateway and create the required API resources and method
P. Use an Amazon API Gateway custom authorizer to control access to the AP
Q. Configure the methods to use AWS Lambda custom integrations, and process each resource with a unique Lambda functio
R. Store request meta information in an Amazon ElastiCache Multi-AZ cluster and static content in a secured S3 bucke
S. Generate presigned URLs when returning references to content stored in Amazon S3.

Answer: C

NEW QUESTION 3
A company is designing a new highly available web application on AWS. The application requires consistent and reliable connectivity from the application servers in AWS to a backend REST API hosted in the company's on-premises environment. The backend connection between AWS and on-premises will be routed over an AWS Direct Connect connection through a private virtual interface. Amazon Route 53 will be used to manage private DNS records for the application to resolve the IP address on the backend REST API.
Which design would provide a reliable connection to the backend API?

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

Answer: A

NEW QUESTION 4
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 runs an application on a fleet of Amazon EC2 instances The application requires low latency and random access to 100 GB of data The application must be able to access the data at up to 3.000 IOPS A Development team has configured the EC2 launch template to provision a 100-GB Provisioned IOPS (PIOPS) Amazon EBS volume with 3 000 IOPS provisioned A Solutions Architect is tasked with lowering costs without impacting performance and durability Which action should be taken?

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

Answer: A

NEW QUESTION 6
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 Consol
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 7
A company runs its containerized batch jobs on Amazon ECS. The jobs are scheduled by submitting a container image, a task definition, and the relevant data to an Amazon S3 bucket. Container images may be unique per job. Running the jobs as quickly as possible is of utmost importance, so submitting jobs artifacts to the S3 bucket triggers the job to run immediately. Sometimes there may no jobs running at all. However, jobs of any size can be submitted with no prior warning to the IT Operations team. Job definitions include CPU and memory resource requirements.
What solution will allow the batch jobs to complete as quickly as possible after being scheduled?

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

Answer: C

NEW QUESTION 8
A Solutions Architect is responsible for redesigning a legacy Java application to improve its availability, data durability, and scalability. Currently, the application runs on a single high-memory Amazon EC2 instance. It accepts HTTP requests from upstream clients, adds them to an in-memory queue, and responds with a 200 status. A separate application thread reads items from the queue, processes them, and persists the results to an Amazon RDS MySQL instance. The processing time for each item takes 90 seconds on average., most of which is spent waiting on external service calls, but the application is written to process multiple items in parallel.
Traffic to this service is unpredictable. During periods of high load, items may sit in the internal queue for over an hour while the application processes the backing. In addition, the current system has issues with availability and data if the single application node fails.
Clients that access this service cannot be modified. They expect to receive a response to each HTTP request they send within 10 seconds before they will time out and retry the request.
Which approach would improve the availability and durability of the system while decreasing the processing latency and minimizing costs?

A. Create an Amazon API Gateway REST API that uses Lambda proxy integration to pass requests to an AWS Lambda functio
B. Migrate the core processing code to a Lambda function and write a wrapper class that provides a handler method that converts the proxy events to the internal application data model and invokes the processing module.
C. Create an Amazon API Gateway REST API that uses a service proxy to put items in an Amazon SQS queu
D. Extract the core processing code from the existing application and update it to pull items from Amazon SQS queu
E. Extract the core processing code from the existing application and update it to pull items from Amazon SQS instead of an in-memory queu
F. Deploy the new processing application to smaller EC2 instances within an Auto Scaling group that scales dynamically based on the approximate number of messages in the Amazon SQS queue.
G. Modify the application to use Amazon DynamoDB instead of Amazon RD

H. Configure Auto Scaling for the DynamoDB tabl
I. Deploy the application within an Auto Scaling group with a scaling policy based on CPU utilizatio
J. Back the in-memory queue with a memory-mapped file to an instance store volume and periodically write that file to Amazon S3.
K. Update the application to use a Redis task queue instead of the in-memory queu
L. Build a Docker container image for the applicatio
M. Create an Amazon ECS task definition that includes the application container and a separate container to host Redi
N. Deploy the new task definition as an ECS service using AWS Fargate and enable Auto Scaling.

Answer: B

NEW QUESTION 9
A company uses Amazon S3 to store documents that may only be accessible to an Amazon EC2 instance in a certain virtual private cloud (VPC). The company fears that a malicious insider with access to this instance could also set up an EC2 instance in another VPC to access these documents.
Which of the following solutions will provide the required protection?

A. Use an S3 VPC endpoint and an S3 bucket policy to limit access to this VPC endpoint.
B. Use EC2 instance profiles and an S3 bucket policy to limit access to the role attached to the instance profile.
C. Use S3 client-side encryption and store the key in the instance metadata.
D. Use S3 server-side encryption and protect the key with an encryption context.

Answer: A

Explanation:
https://docs.aws.amazon.com/vpc/latest/userguide/vpce-gateway.html
Endpoint connections cannot be extended out of a VPC. Resources on the other side of a VPN connection, VPC peering connection, AWS Direct Connect connection, or ClassicLink connection in your VPC cannot use the endpoint to communicate with resources in the endpoint service.

NEW QUESTION 10
A company is operating a large customer service call center, and stores and processes call recordings with a custom application Approximately 2% of the call recording are transcribed by an offshore team for quality assurance purposes. These recordings take days. The company uses Linux servers for processing the call recording and managing the transcription queue. There is also a web application for the quality assurance staff to review and score call recordings.
The company plans to migrate the system to AWS to reduce storage costs and the time required to transcribe calls.
Which set of actions should be taken to meet the company's objectives?

A. Upload the call recording to Amazon S3 from the call cente
B. Set up an S3 lifecycle policy to move the call recordings to Amazon S3 Glacier after 90 day
C. Use an AWS Lambda trigger to transcribe the call recordings with Amazon Transcrib
D. Use Amazon S3, Amazon API Gateway and Lambda to host the review and scoring application.
E. Upload the call recordings to Amazon S3 from the call cente
F. Set up an S3 lifecycle policy to move the call recordings to Amazon S3 Glacier after 90 day
G. Use an AWS Lambda trigger to transcribe the call recordings with Amazon Mechanical trun
H. Use Amazon EC2 instances in an Auto Scaling group behind an Application Balancer to host the review and scoring application.
I. Use Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer to host the review and scoring applicatio
J. Upload the call recordings to this application from the call center and store them on an Amazon EFS mount poin
K. Use AWS Backup to archive the call recording after 90 day
L. Transcribe the call recordings with Amazon Transcribe.
M. Upload the call recording to Amazon S3 from the call center and put the object key in an Amazon SQS 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 10
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 12
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 14
A development team has created a series of AWS CloudFormation templates to help deploy services. They created a template for a network/virtual private (VPC) stack, a database stack, a bastion host stack, and a web application-specific stack. Each service requires the deployment of at least:
Each template has multiple input parameters that make it difficult to deploy the services individually from the AWS CloudFormation console. The input parameters from one stack are typically outputs from other stacks. For example, the VPC ID, subnet IDs, and security groups from the network stack may need to be used in the application stack or database stack.
Which actions will help reduce the operational burden and the number of parameters passed into a service deployment? (Choose two.)

A. Create a new AWS CloudFormation template for each servic
B. After the existing templates to use cross-stack references to eliminate passing many parameters to each templat
C. Call each required stack for the application as a nested stack from the new stac
D. Call the newly created service stack from theAWS CloudFormation console to deploy the specific service with a subset of the parameters previously required.
E. Create a new portfolio in AWS Service Catalog for each servic
F. Create a product for each existing AWS CloudFormation template required to build the servic
G. Add the products to the portfolio that represents that service in AWS Service Catalo
H. To deploy the service, select the specific service portfolio and launch the portfolio with the necessary parameters to deploy all templates.
I. Set up an AWS CodePipeline workflow for each servic
J. For each existing template, choose AWS CloudFormation as a deployment actio
K. Add the AWS CloudFormation template to the deployment actio
L. Ensure that the deployment actions are processed to make sure that dependences are obeye
M. Use configuration files and scripts to share parameters between the stack
N. To launch the service, execute the specific template by choosing the name of the service and releasing a change.
O. Use AWS Step Functions to define a new servic
P. Create a new AWS CloudFormation template for each servic
Q. After the existing templates to use cross-stack references to eliminate passing many parameters to each templat
R. Call each required stack for the application as a nested stack from the new service templat
S. Configure AWS Step Functions to call the service template directl
T. In the AWS Step Functions console, execute the step.
. Create a new portfolio for the Services in AWS Service Catalo
. Create a new AWS CloudFormation template for each servic
. After the existing templates to use cross-stack references to eliminate passing many parameters to each templat
. Call each required stack for the application as a nested stack from the new stac
. Create a product for each applicatio
. Add the service template to the produc
. Add each new product to the portfoli
. Deploy the product from the portfolio to deploy the service with the necessary parameters only to start the deployment.

Answer: AE

NEW QUESTION 18
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 22
A company has an Amazon VPC that is divided into a public subnet and a private subnet A web application runs in Amazon VPC, and each subnet has its own NACL The public subnet has a CIDR of 10.0.0.0/24. An Application Load Balancer is deployed to the public subnet. The private subnet has a CIDR of 10.0.1.0/24.
Amazon EC2 instances that run a web server on port 80 are launched into the private subnet.
Only network traffic that is required for the Application Load Balancer to access the web application can be allowed to travel between the public and private subnets
What collection of rules should be written to ensure that the private subnet's NACL meets the requirement? (Select TWO.)

A. An inbound rule for port 80 from source 0.0.0 0/0
B. An inbound rule for port 80 from source 10.0.0.0/24
C. An outbound rule for port 80 to destination 0.0.0.0/0
D. An outbound rule for port 80 to destination 10.0.0.0/24
E. An outbound rule for ports 1024 through 65535 to destination 10.0.0.0/24

Answer: BE

NEW QUESTION 26
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 29
A company wants to ensure that the workloads for each of its business units have complete autonomy and a minimal blast radius in AWS. The Security team must be able to control access to the resources and services in the account to ensure that particular services are not used by the business units.
How can a Solutions Architect achieve the isolation requirements?

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

Answer: A

NEW QUESTION 33
A company is using AWS for production and development workloads. Each business unit has its own AWS account for production, and a separate AWS account to develop and deploy its applications. The Information Security department has introduced new security policies that limit access for terminating certain Amazon ECs instances in all accounts to a small group of individuals from the Security team.
How can the Solutions Architect meet these requirements?

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

Answer: B

NEW QUESTION 36
A company's CISO has asked a Solutions Architect to re-engineer the company's current CI/CD practices to make sure patch deployments to its applications can happen as quickly as possible with minimal downtime if vulnerabilities are discovered. The company must also be able to quickly roll back a change in case of errors. The web application is deployed in a fleet of Amazon EC2 instances behind an Application Load Balancer. The company is currently using GitHub to host the application source code and has configured an AWS CodeBuild project to build the application. The company also intends to use AWS CodePipeLine to trigger builds form GitHub commits using the existing CodeBuild project.
What CI/CD configuration meets all of the requirements?

A. Configure CodePipeline with a deploy stage using AWS CodeDeploy configured for in-place deploymen
B. Monitor the newly deployed code, and if there are any issues, push another code update.
C. Configure CodePipeline with a deploy stage using AWS CodeDeploy configure for blue/green deployment
D. Monitor the new deployed code and if there are any issues, trigger a manual rollback using CodeDeploy.
E. Configure CodePipeline with a deploy stage using AWS CloudFormation to create a pipeline for test and production stack
F. Monitor the newly deployed cod and if there are any issues push another code update.
G. Configure the CodePipeline with a deploy stage using AWS OpsWorks and in-place deployments.Monitor the newly deployed code and if there are any issues, push another code update.

Answer: B

NEW QUESTION 38
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 41
A company used Amazon EC2 instances to deploy a web fleet to host a blog site. The EC2 instances are behind an Application Load Balancer (ALB) and are configured in an Auto Scaling group. The web application stores all blog content on an Amazon EFS volume.
The company recently added a feature for bloggers to add video to their posts, attracting 10 times the previous user traffic. At peak times of day, users report buffering and timeout issues while attempting to reach the site or watch videos.
Which is the MOST cost-efficient and scalable deployment that will resolve the issues for users?

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

Answer: C

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

NEW QUESTION 42
A Solutions Architect must update an application environment within AWS Elastic Beanstalk using a blue/green deployment methodology. The Solutions Architect creates an environment that is identical to the existing application environment and deploys the application to the new environment.
What should be done next to complete the update?

A. Redirect to the new environment using Amazon Route 53
B. Select the Swap Environment URLs option
C. Replace the Auto Scaling launch configuration
D. Update the DNS records to point to the green environment

Answer: B

Explanation:
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.CNAMESwap.html

NEW QUESTION 47
A group of research institutions and hospitals are in a partnership to study 2 PBs of genomic data. The institute that owns the data stores it in an Amazon S3 bucket and updates it regularly. The institute would like to give all of the organizations in the partnership read access to the data. All members of the partnership are extremely cost-conscious, and the institute that owns the account with the S3 bucket is concerned about covering the costs for requests and data transfers from Amazon S3.
Which solution allows for secure datasharing without causing the institute that owns the bucket to assume all the costs for S3 requests and data transfers?

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

Answer: B

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

NEW QUESTION 51
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 54
A company is creating an account strategy so that they can begin using AWS. The Security team will provide each team with the permissions they need to follow the principle or least privileged access. Teams would like to keep their resources isolated from other groups, and the Finance team would like each team's resource usage separated for billing purposes.
Which account creation process meets these requirements and allows for changes?

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

Answer: B

NEW QUESTION 58
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 queu
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 bucke
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 bucke
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 bucke
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 bucke
J. Use the queue depth metric to adjust the size of the Auto Scaling group for video processing instances.

Answer: A

NEW QUESTION 59
A company has been using a third-party provider for its content delivery network and recently decided to switch to Amazon CloudFront the Development team wants to maximize performance for the global user base. The company uses a content management system (CMS) that serves both static and dynamic content. The CMS is both md an Application Load Balancer (ALB) which is set as the default origin for the distribution. Static assets are served from an Amazon S3 bucket. The Origin Access Identity (OAI) was created property d the S3 bucket policy has been updated to allow the GetObject action from the OAI, but static assets are receiving a 404 error
Which combination of steps should the Solutions Architect take to fix the error? (Select TWO.)

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

Answer: AD

NEW QUESTION 61
A company has a legacy application running on servers on premises. To increase the application's reliability, the company wants to gain actionable insights using application logs. A Solutions Architect has been given following requirements for the solution:

> Aggregate logs using AWS.

> Automate log analysis for errors.

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

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

Answer: A

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

NEW QUESTION 65
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 66
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 68
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 applianc
B. Create a VPN connection to each VP
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 gatewa
E. Default route internet traffic back to an on-premises router to route to the internet.
F. Create a central VPC for outbound internet traffi
G. Use VPC peering to default route to a set of redundant NAT gateway in the central VPC.
H. Create a proxy fleet in a central VPC accoun
I. Create an AWS PrivateLink endpoint service in the central VP
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 71
An e-commerce company is revamping its IT infrastructure and is planning to use AWS services. The company's CIO has asked a Solutions Architect to design a simple, highly available, and loosely coupled order processing application. The application is responsible for receiving and processing orders before storing them in an Amazon DynamoDB table. The application has a sporadic traffic pattern and should be able to scale during marketing campaigns to process the orders with minimal delays.
Which of the following is the MOST reliable approach to meet the requirements?

A. Receive the orders in an Amazon EC2-hosted database and use EC2 instances to process them.
B. Receive the orders in an Amazon SQS queue and trigger an AWS Lambda function to process them.
C. Receive the orders using the AWS Step Functions program and trigger an Amazon ECS container to process them.
D. Receive the orders in Amazon Kinesis Data Streams and use Amazon EC2 instances to process them.

Answer: B

NEW QUESTION 72
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 73
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 78
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 79
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 deleteon 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 toeu-west-1.

Answer: A

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

NEW QUESTION 83
A Solutions Architect is designing a multi-account structure that has 10 existing accounts. The design must meet the following requirements:

> Consolidate all accounts into one organization.

> Allow full access to the Amazon EC2 service from the master account and the secondary accounts.

> Minimize the effort required to add additional secondary accounts.
Which combination of steps should be included in the solution? (Choose two.)

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

Answer: AD

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

NEW QUESTION 87
A company is finalizing the architecture for its backup solution for applications running on AWS. All of the applications run on AWS and use at least two Availability Zones in each tier.
Company policy requires IT to durably store nightly backups for all its data in at least two locations: production and disaster recovery. The locations must be in different geographic regions. The company also needs the backup to be available to restore immediately at the production data center, and within 24 hours at the disaster recovery location. All backup processes must be fully automated.
What is the MOST cost-effective backup solution that will meet all requirements?

A. Back up all the data to a large Amazon EBS volume attached to the backup media server in the production regio
B. Run automated scripts to snapshot these volumes nightly, and copy these snapshots to the disaster recovery region.
C. Back up all the data to Amazon S3 in the disaster recovery regio
D. Use a lifecycle policy to move this data to Amazon Glacier in the production region immediatel
E. Only the data is replicated; remove the data from the S3 bucket in the disaster recovery region.
F. Back up all the data to Amazon Glacier in the production regio
G. Set up cross-region replication of this data to Amazon Glacier in the disaster recovery regio
H. Set up a lifecycle policy to delete any data older than 60 days.
I. Back up all the data to Amazon S3 in the production regio
J. Set up cross-region replication of this S3 bucket to another region and set up a lifecycle policy in the second region to immediately move this data to Amazon Glacier.

Answer: D

NEW QUESTION 91
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 92
A company has a web application that securely uploads pictures and videos to an Amazon S3 bucket The company requires that only authenticated users are allowed to post content The application generates a presigned URL that is used to upload objects through a browser interface Most users are reporting slow upload times for objects larger than 100 MB.
What can a Solutions Architect do to improve the performance of these uploads while ensuring only authenticated users are allowed to post content?

A. Set up an Amazon API Gateway with an edge-optimized API endpoint that has a resource as an S3 service proxy Configure the PUT method for this resource to expose the S3 Putobject operation Secure the API Gateway using a COGNITO_USER_POOLS authorize
B. Have the browser interface use API Gateway instead of the presigned URL to upload objects
C. Set up an Amazon API Gateway with a regional API endpoint that has a resource as an S3 service proxy Configure the PUT method for this resource to expose the S3 Putobject operation Secure the API Gateway using an AWS Lambda authorizer Have the browser interface use API Gateway instead of the presigned URL lo upload objects
D. Enable an S3 Transfer Acceleration endpoint on the S3 bucket Use the endpoint when generating the presigned URL Have the browser interface upload the objects to the URL using the S3 multipart upload API.
E. Configure an Amazon CloudFront distribution for the destination S3 bucket Enable PUT and POST methods for the CloudFront cache behavior Update the CloudFront origin to use an origin access identity (OAI). Give the OAI user s3:PutObject permissions in the bucket policy Have the browser interface upload objects using the CloudFront distribution.

Answer: A

NEW QUESTION 96
A large global company wants to migrate a stateless mission-critical application to AWS. The application is based on IBM WebSphere (application and integration middleware), IBM MQ (messaging middleware), and IBM DB2 (database software) on a z/OS operating system.
How should the Solutions Architect migrate the application to AWS?

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

Answer: B

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

NEW QUESTION 100
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 103
A company has released a new version of a website to target an audience in Asia and South America. The website's media assets are hosted on Amazon S3 and have an Amazon CloudFront distribution to improve end-user performance. However, users are having a poor login experience the authentication service is only available in the us-east-1 AWS Region.
How can the Solutions Architect improve the login experience and maintain high security and performance with minimal management overhead?

A. Replicate the setup in each new geography and use Amazon Route 53 geo-based routing to route traffic to the AWS Region closest to the users.
B. Use an Amazon Route 53 weighted routing policy to route traffic to the CloudFront distributio
C. Use CloudFront cached HTTP methods to improve the user login experience.
D. Use Amazon Lambda@Edge attached to the CloudFront viewer request trigger to authenticate and authorize users by maintaining a secure cookie token with a session expiry to improve the user experience in multiple geographies.
E. Replicate the setup in each geography and use Network Load Balancers to route traffic to the authentication service running in the closest region to users.

Answer: C

Explanation:
There are several benefits to using Lambda@Edge for authorization operations. First, performance is improved by running the authorization function using Lambda@Edge closest to the viewer, reducing latency and response time to the viewer request. The load on your origin servers is also reduced by offloading CPU-intensive operations such as verification of JSON Web Token (JWT) signatures. Finally, there are security benefits such as filtering out unauthorized requests before they reach your origin infrastructure.
https://aws.amazon.com/blogs/networking-and-content-delivery/authorizationedge-how-to-use-lambdaedge-and-

NEW QUESTION 104
A company prefers to limit running Amazon EC2 instances to those that were launched from AMIs pre-approved by the Information Security department. The Development team has an agile continuous integration and deployment process that cannot be stalled by the solution.
Which method enforces the required controls with the LEAST impact on the development process? (Choose two.)

A. Use IAM policies to restrict the ability of users or other automated entities to launch EC2 instances based on a specific set of pre-approved AMIs, such as those tagged in a specific way by Information Security.
B. Use regular scans within Amazon Inspector with a custom assessment template to determine if the EC2 instance that the Amazon Inspector Agent is running on is based upon a pre-approved AM
C. If it is not, shut down the instance and inform information Security by email that this occurred.
D. Only allow launching of EC2 instances using a centralized DevOps team, which is given work packages via notifications from an internal ticketing syste
E. Users make requests for resources using this ticketing tool, which has manual information security approval steps to ensure that EC2 instances are only launched from approved AMIs.
F. Use AWS Config rules to spot any launches of EC2 instances based on non-approved AMIs, trigger an AWS Lambda function to automatically terminate the instance, and publish a message to an Amazon SNS topic to inform Information Security that this occurred.
G. Use a scheduled AWS Lambda function to scan through the list of running instances within the virtual private cloud (VPC) and determine if any of these are based on unapproved AMI
H. Publish a message to an SNS topic to inform Information Security that this occurred and then shut down the instance.

Answer: AD

Explanation:
https://docs.aws.amazon.com/config/latest/developerguide/evaluate-config_develop-rules_getting-started.html

NEW QUESTION 106

......

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