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 company wants to follow its website on AWS using serverless architecture design patterns for global customers. The company has outlined its requirements as follow:
> The website should be responsive.
> The website should offer minimal latency.
> The website should be highly available.
> Users should be able to authenticate through social identity providers such as Google, Facebook, and Amazon.
> There should be baseline DDoS protections for spikes in traffic.
How can the design requirements be met?

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

Answer: C

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

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

Answer: A

NEW QUESTION 4
A photo-sharing and publishing company receives 10,000 to 150,000 images daily. The company receives the images from multiple suppliers and users registered with the service. The company is moving to AWS and wants to enrich the existing metadata by adding data using Amazon Rekognition.
The following is an example of the additional data:

```
list celebrities [name of the personality] wearing [color] looking [happy, sad] near [location example Eiffel Tower in Paris]
```

As part of the cloud migration program, the company uploaded existing image data to Amazon S3 and told users to upload images directly to Amazon S3.
What should the Solutions Architect do to support these requirements?

A. Trigger AWS Lambda based on an S3 event notification to create additional metadata using Amazon Rekognitio
B. Use Amazon DynamoDB to store the metadata and Amazon ES to create an inde
C. Use a web front-end to provide search capabilities backed by Amazon ES.
D. Use Amazon Kinesis to stream data based on an S3 even
E. Use an application running in Amazon EC2 to extract metadata from the image
F. Then store the data on Amazon DynamoDB and Amazon CloudSearch and create an inde
G. Use a web front-end with search capabilities backed by CloudSearch.
H. Start an Amazon SQS queue based on S3 event notification

I. Then have Amazon SQS send the metadata information to Amazon DynamoD
J. An application running on Amazon EC2 extracts data from Amazon Rekognition using the API and adds data to DynamoDB and Amazon E
K. Use a web front-end to provide search capabilities backed by Amazon ES.
L. Trigger AWS Lambda based on an S3 event notification to create additional metadata using Amazon Rekognitio
M. Use Amazon RDS MySQL Multi-AZ to store the metadata information and use Lambda to create an inde
N. Use a web front-end with search capabilities backed by Lambda.

Answer: A

Explanation:
https://github.com/aws-samples/lambda-refarch-imagerecognition

NEW QUESTION 5
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 6
A company CFO recently analyzed the company's AWS monthly bill and identified an opportunity to reduce the cost for AWS Elastic Beanstalk environments in use. The CFO has asked a Solutions Architect to design a highly available solution that will spin up an Elastic Beanstalk environment in the morning and terminate it at the end of the day.
The solution should be designed with minimal operational overhead and to minimize costs. It should also be able to handle the increased use of Elastic Beanstalk environments among different teams, and must provide a one-stop scheduler solution for all teams to keep the operational costs low.
What design will meet these requirements?

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

Answer: B

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

NEW QUESTION 7
A company is using AWS CloudFormation to deploy its infrastructure. The company is concerned that, if a production CloudFormation stack is deleted, important data stored in Amazon RDS databases or Amazon EBS volumes might also be deleted.

How can the company prevent users from accidentally deleting data in this way?

A. Modify the CloudFormation templates to add a DeletionPolicy attribute to RDS and EBS resources.
B. Configure a stack policy that disallows the deletion of RDS and EBS resources.
C. Modify IAM policies to deny deleting RDS and EBS resources that are tagged with an "aws:cloudformation:stack-name" tag.
D. Use AWS Config rules to prevent deleting RDS and EBS resources.

Answer: A

Explanation:
With the DeletionPolicy attribute you can preserve or (in some cases) backup a resource when its stack is deleted. You specify a DeletionPolicy attribute for each resource that you want to control. If a resource has no DeletionPolicy attribute, AWS CloudFormation deletes the resource by default. To keep a resource when its stack is deleted, specify Retain for that resource. You can use retain for any resource. For example, you can retain a nested stack, Amazon S3 bucket, or EC2 instance so that you can continue to use or modify those resources after you delete their stacks.
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-deletionpolicy.html

NEW QUESTION 8
A company has developed a new billing application that will be released in two weeks. Developers are testing the application running on 10 EC2 instances managed by an Auto Scaling group in subnet 172.31.0.0/24 within VPC A with CIDR block 172.31.0.0/16. The Developers noticed connection timeout errors in the application logs while connecting to an Oracle database running on an Amazon EC2 instance in the same region within VPC B with CIDR block 172.50.0.0/16. The IP of the database instance is hard-coded in the application instances.
Which recommendations should a Solutions Architect present to the Developers to solve the problem in a secure way with minimal maintenance and overhead?

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

Answer: C

NEW QUESTION 9
A company is migrating a subset of its application APIs from Amazon EC2 instances to run on a serverless infrastructure. The company has set up Amazon API Gateway, AWS Lambda, and Amazon DynamoDB for the new application. The primary responsibility of the Lambda function is to obtain data from a third-party Software as a Service (SaaS) provider. For consistency, the Lambda function is attached to the same virtual private cloud (VPC) as the original EC2 instances. Test users report an inability to use this newly moved functionality, and the company is receiving 5xx errors from API Gateway. Monitoring reports from the SaaS provider shows that the requests never made it to its systems. The company notices that Amazon CloudWatch Logs are being generated by the Lambda functions. When the same functionality is tested against the EC2 systems, it works as expected.
What is causing the issue?

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

Answer: A

NEW QUESTION 10
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 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 15
A bank is designing an online customer service portal where customers can chat with customer service agents. The portal is required to maintain a 15-minute RPO or RTO in case of a regional disaster. Banking regulations require that all customer service chat transcripts must be preserved on durable storage for at least 7 years, chat conversations must be encrypted in-flight, and transcripts must be encrypted at rest. The Data Lost Prevention team requires that data at rest must be encrypted using a key that the team controls, rotates, and revokes.
Which design meets these requirements?

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

Answer: B

NEW QUESTION 17
A Solutions Architect wants to make sure that only AWS users or roles with suitable permissions can access a new Amazon API Gateway endpoint The Solutions Architect wants an end-to-end view of each request to analyze the latency of the request and create service maps
How can the Solutions Architect design the API Gateway access control and perform request inspections?

A. For the API Gateway method set the authorization to AWSJAM Then, give the I AM user or role execute-api Invoke permission on the REST API resource Enable the API caller to sign requests with AWS Signature when accessing the endpoint Use AWS X-Roy to trace and analyze user requests to API Gateway
B. For the API Gateway resource set CORS to enabled and only return the company's domain mAccess-Control-Allow-Origin headers Then give the IAM user or role execute-api Invoke permission on the REST API resource Use Amazon CloudWatch to trace and analyze user requests to API Gateway
C. Create an AWS Lambda function as the custom authorizer ask the API client to pass the key and secret when making the call and then use Lambda to validate the key'secret pair against the IAM system Use AWS X-Ray to trace and analyze user requests to API Gateway
D. Create a client certificate for API Gateway Distribute the certificate to the AWS users and roles that need to access the endpoint Enable the API caller to pass the client certificate when accessing the endpoint Use Amazon CloudWatch to trace and analyze user requests to API Gateway.

Answer: D

NEW QUESTION 21
A company runs a public-facing application that uses a Java-based web sen/ice via a RESTful API It is hosted on Apache Tomcat on a single server in a data center that runs consistently at 30% CPU utilization Use of the API is expected to increase by 10 times with a new product launch The business wants to migrate the application to AWS with no disruption and needs it to scale to meet demand
The company has already decided to use Amazon Route 53 and CNAME records lo redirect traffic How can these requirements be met with the LEAST amount of effort?

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

Answer: A

NEW QUESTION 22
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 24
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 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 30
A company stores sales transaction data in Amazon DynamoDB tables. To detect anomalous behaviors and respond quickly, all changes to the items stored in the DynamoDB tables must be logged within 30 minutes. Which solution meets the requirements?

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

Answer: C

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

NEW QUESTION 32
A company runs a dynamic mission-critical web application that has an SLA of 99.99%. Global application users access the application 24/7. The application is currently hosted on premises and routinely fails to meet its SLA, especially when millions of users access the application concurrently. Remote users complain of latency.
How should this application be redesigned to be scalable and allow for automatic failover at the lowest cost?

A. Use Amazon Route 53 failover routing with geolocation-based 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 34
A company wants to manage the costs associated with a group of 20 applications that are critical, by migrating to AWS. The applications are a mix of Java and Node.js spread across different instance clusters. The company wants to minimize costs while standardizing by using a single deployment methodology. Most of the applications are part of month-end processing routines with a small number of concurrent users, but they are occasionally run at other times. Average application memory consumption is less than 1 GB, though some applications use as much as 2.5 GB of memory during peak processing. The most important application in the group is a billing report written in Java that accesses multiple data sources and often for several hours.
Which is the MOST cost-effective solution?

A. Deploy a separate AWS Lambda function for each applicatio
B. Use AWS CloudTrail logs and Amazon CloudWatch alarms to verify completion of critical jobs.
C. Deploy Amazon ECS containers on Amazon EC2 with Auto Scaling configured for memory utilization of 75%. Deploy an ECS task for each application being migrated with ECS task scalin
D. Monitor services and hosts by using Amazon CloudWatch.
E. Deploy AWS Elastic Beanstalk for each application with Auto Scaling to ensure that all requests have sufficient resource
F. Monitor each AWS Elastic Beanstalk deployment with using CloudWatch alarms.
G. Deploy a new amazon EC2 instance cluster that co-hosts all applications by using EC2 Auto Scaling and Application Load Balancer
H. Scale cluster size based on a custom metric set on instance memory utilizatio
I. Purchase 3-year Reserved instance reservations equal to the GroupMaxSize parameter of the Auto Scaling group.

Answer: C

NEW QUESTION 39
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 43
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 45
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 50
A Solutions Architect is designing the storage layer for a recently purchased application. The application will be running on Amazon EC2 instances and has the following layers and requirements:

> Data layer: A POSIX file system shared across many systems.

> Service layer: Static file content that requires block storage with more than 100k IOPS. Which combination of AWS services will meet these needs? (Choose two.)

A. Data layer – Amazon S3
B. Data layer – Amazon EC2 Ephemeral Storage
C. Data layer – Amazon EFS
D. Service layer – Amazon EBS volumes with Provisioned IOPS
E. Service layer – Amazon EC2 Ephemeral Storage

Answer: CE

Explanation:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/storage-optimized-instances.html

NEW QUESTION 54
A company runs a legacy system on a single m4.2xlarge Amazon EC2 instance with Amazon EBS2 storage. The EC2 instance runs both the web server and a self-managed Oracle database. A snapshot is made of the EBS volume every 12 hours, and an AMI was created from the fully configured EC2 instance.
A recent event that terminated the EC2 instance led to several hours of downtime. The application was successfully launched from the AMI, but the age of the EBS snapshot and the repair of the database resulted in the loss of 8 hours of data. The system was also down for 4 hours while the Systems Operators manually performed these processes.
What architectural changes will minimize downtime and reduce the chance of lost data?

A. Create an Amazon CloudWatch alarm to automatically recover the instanc
B. Create a script that will check and repair the database upon reboo
C. Subscribe the Operations team to the Amazon SNS message generated by the CloudWatch alarm.
D. Run the application on m4.xlarge EC2 instances behind an Elastic Load Balancer/Application Load Balance
E. Run the EC2 instances in an Auto Scaling group across multiple Availability Zones with a minimum instance count of tw
F. Migrate the database to an Amazon RDS Oracle Multi-AZ DB instance.
G. Run the application on m4.2xlarge EC2 instances behind an Elastic Load Balancer/Application Load Balance
H. Run the EC2 instances in an Auto Scaling group across multiple Availability Zones with aminimum instance count of on
I. Migrate the database to an Amazon RDS Oracle Multi-AZ DB instance.

J. Increase the web server instance count to two m4.xlarge instances and use Amazon Route 53 round-robin load balancing to spread the loa
K. Enable Route 53 health checks on the web server
L. Migrate the database to an Amazon RDS Oracle Multi-AZ DB instance.

Answer: B

Explanation:
Ensures that there are at least two EC instances, each of which is in a different AZ. It also ensures that the database spans multiple AZs. Hence this meets all the criteria.
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html

NEW QUESTION 55
What combination of steps could a Solutions Architect take to protect a web workload running on Amazon EC2 from DDoS and application layer attacks? (Select two.)

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

Answer: BD

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

NEW QUESTION 57
A media company has a 30-TB repository of digital news videos. These videos are stored on tape in an
on-premises tape library and referenced by a Media Asset Management (MAM) system. The company wants to enrich the metadata for these videos in an
automated fashion and put them into a searchable catalog by using a MAM feature. The company must be able to search based on information in the video, such
as objects, scenery items, or people's faces. A catalog is available that contains faces of people who have appeared in the videos that include an image of each
person. The company would like to migrate these videos to AWS.
The company has a high-speed AWS Direct Connect connection with AWS and would like to move the MAM solution video content directly from its current file
system.
How can these requirements be met by using the LEAST amount of ongoing management overhead and causing MINIMAL disruption to the existing system?

A. Set up an AWS Storage Gateway, file gateway appliance on-premise
B. Use the MAM solution to extract the videos from the current archive and push them into the file gatewa
C. Use the catalog of faces to build a collection in Amazon Rekognitio
D. Build an AWS Lambda function that invokes the Rekognition Javascript SDK to have Rekognition pull the video from the Amazon S3 files backing the file
gateway, retrieve the required metadata, and push the metadata into the MAM solution.
E. Set up an AWS Storage Gateway, tape gateway appliance on-premise
F. Use the MAM solution to extract the videos from the current archive and push them into the tape gatewa
G. Use the catalog of faces to build a collection in Amazon Rekognitio
H. Build an AWS Lambda function that invokes the Rekognition Javascript SDK to have Amazon Rekognition process the video in the tape gateway, retrieve the
required metadata, and push the metadata into the MAM solution.
I. Configure a video ingestion stream by using Amazon Kinesis Video Stream
J. Use the catalog of faces to build a collection in Amazon Rekognitio
K. Stream the videos from the MAM solution into Kinesis Video Stream
L. Configure Amazon Rekognition to process the streamed video
M. Then, use a stream consumer to retrieve the required metadata, and push the metadata into the MAM solutio
N. Configure the stream to store the videos in Amazon S3.
O. Set up an Amazon EC2 instance that runs the OpenCV librarie
P. Copy the videos, images, and face catalog from the on-premises library into an Amazon EBS volume mounted on this EC2 instanc
Q. Process the videos to retrieve the required metadata, and push the metadata into the MAM solution while also copying the video files to an Amazon S3 bucket.

Answer: C

Explanation:
https://docs.aws.amazon.com/rekognition/latest/dg/streaming-video.html

NEW QUESTION 62
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 65
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 66
A Solutions Architect has been asked to look at a company's Amazon Redshift cluster, which has quickly become an integral part of its technology and supports key business process. The Solutions Architect is to increase the reliability and availability of the cluster and provide options to ensure that if an issue arises, the cluster can either operate or be restored within four hours.
Which of the following solution options BEST addresses the business need in the most cost-effective manner?

A. Ensure that the Amazon Redshift cluster has been set up to make use of Auto Scaling groups with the nodes in the cluster spread across multiple Availability Zones.
B. Ensure that the Amazon Redshift cluster creation has been template using AWS CloudFormation so it can easily be launched in another Availability Zone and data populated from the automated Redshift back-ups stored in Amazon S3.
C. Use Amazon Kinesis Data Firehose to collect the data ahead of ingestion into Amazon Redshift and create clusters using AWS CloudFormation in another region and stream the data to both clusters.
D. Create two identical Amazon Redshift clusters in different regions (one as the primary, one as the secondary). Use Amazon S3 cross-region replication from the primary to secondary). Use Amazon S3 cross-region replication from the primary to secondary region, which triggers an AWS Lambda function to populate the cluster in the secondary region.

Answer: B

Explanation:
https://aws.amazon.com/redshift/faqs/?nc1=h_ls Q: What happens to my data warehouse cluster availability and data durability if my data warehouse cluster's Availability Zone (AZ) has an outage? If your Amazon Redshift data warehouse cluster's Availability Zone becomes unavailable, you will not be able to use your cluster until power and network access to the AZ are restored. Your data warehouse cluster's data is preserved so you can start using your Amazon Redshift data warehouse as soon as the AZ becomes available again. In addition, you can also choose to restore any existing snapshots to a new AZ in the same Region. Amazon Redshift will restore your most frequently accessed data first so you can resume queries as quickly as possible.
FROM 37

NEW QUESTION 71
A Solutions Architect is migrating a 10 TB PostgreSQL database to Amazon RDS for PostgreSQL. The company's internet link is 50 MB with a VPN in the Amazon VPC, and the Solutions Architect needs to migrate the data and synchronize the changes before the cutover. The cutover must take place within an 8-day period.
What is the LEAST complex method of migrating the database securely and reliably?

A. Order an AWS Snowball device and copy the database using the AWS DM
B. When the database is available in Amazon 3, use AWS DMS to load it to Amazon RDS, and configure a job to synchronize changes before the cutover.
C. Create an AWS DMS job to continuously replicate the data from on premises to AW
D. Cutover to Amazon RDS after the data is synchronized.
E. Order an AWS Snowball device and copy a database dump to the devic
F. After the data has been copied to Amazon S3, import it to the Amazon RDS instanc
G. Set up log shipping over a VPN to synchronize changes before the cutover.
H. Order an AWS Snowball device and copy the database by using the AWS Schema Conversion Tool.When the data is available in Amazon S3, use AWS DMS to load it to Amazon RDS, and configure a job to synchronize changes before the cutover.

Answer: B

NEW QUESTION 73
A company's application is increasingly popular and experiencing latency because of high volume reads on the database server.
The service has the following properties:

> A highly available REST API hosted in one region using Application Load Balancer (ALB) with auto scaling.

> A MySQL database hosted on an Amazon EC2 instance in a single Availability Zone.
The company wants to reduce latency, increase in-region database read performance, and have multi-region disaster recovery capabilities that can perform a live recovery automatically without any data or performance loss (HA/DR).

Which deployment strategy will meet these requirements?

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

Answer: A

NEW QUESTION 77
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 79
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 83
An enterprise company is using a multi-account AWS strategy There are separate accounts tor development staging and production workloads To control costs and improve governance the following requirements have been defined:
• The company must be able to calculate the AWS costs tor each project
• The company must be able to calculate the AWS costs tor each environment development staging and production
• Commonly deployed IT services must be centrally managed
• Business units can deploy pre-approved IT services only

• Usage of AWS resources in the development account must be limited
Which combination of actions should be taken to meet these requirements? (Select THREE)

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

Answer: CEF

NEW QUESTION 84
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 87
A company that provides wireless services needs a solution to store and analyze log files about user activities. Currently, log files are delivered daily to Amazon Linux on Amazon EC2 instance. A batch script is run once a day to aggregate data used for analysis by a third-party tool. The data pushed to the third-party tool is used to generate a visualization for end users. The batch script is cumbersome to maintain, and it takes several hours to deliver the ever-increasing data volumes to the third-party tool. The company wants to lower costs, and is open to considering a new tool that minimizes development effort and lowers administrative overhead. The company wants to build a more agile solution that can store and perform the analysis in near-real time, with minimal overhead. The solution needs to be cost effective and scalable to meet the company's end-user base growth.
Which solution meets the company's requirements?

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

Answer: B

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

NEW QUESTION 90
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 95
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 98
A Solutions Architect is designing the storage layer for a data warehousing application. The data files are large, but they have statically placed metadata at the beginning of each file that describes the size and placement of the file's index. The data files are read in by a fleet of Amazon EC2 instances that store the index size, index location, and other category information about the data file in a database. That database is used by Amazon EMR to group files together for deeper analysis.
What would be the MOST cost-effective, high availability storage solution for this workflow?

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

Answer: A

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

NEW QUESTION 102
A company wants to replace its call system with a solution built using AWS managed services. The company call center would like the solution to receive calls, create contact flows, and scale to handle growth projections. The call center would also like the solution to use deep learning capabilities to recognize the intent of the callers and handle basic tasks, reducing the need to speak an agent. The solution should also be able to query business applications and provide relevant information back to calls as requested.
Which services should the Solution Architect use to build this solution? (Choose three.)

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

Answer: BDE

NEW QUESTION 105
A company has a website that enables users to upload videos Company policy states the uploaded videos must be analyzed for restricted content An uploaded video is placed in Amazon S3, and a message is pushed to an Amazon SQS queue with the video's location A backend application pulls this location from Amazon SQS and analyzes the video
The video analysis is compute-intensive and occurs sporadically during the day The website scales with demand The video analysis application runs on a fixed number of instances Peak demand occurs during the holidays, so the company must add instances to the application during this time All instances used are currently on-demand Amazon EC2 T2 instances The company wants to reduce the cost of the current solution.
Which of the following solutions is MOST cost-effective?

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

Answer: B

NEW QUESTION 109
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 112
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 patche
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 113
A company needs to run a software package that has a license that must be run on the same physical host for the duration of its use. The software package is only going to be used for 90 days. The company requires patching and restarting of all instances every 30 days.
How can these requirements be met using AWS?

A. Run a dedicated instance with auto-placement disabled.
B. Run the instance on a dedicated host with Host Affinity set to Host.
C. Run an On-Demand instance with a Reserved Instance to ensure consistent placement.
D. Run the instance on a licensed host with termination set for 90 days.

Answer: B

Explanation:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/how-dedicated-hosts-work.html

NEW QUESTION 114
A Solutions Architect is working with a company that operates a standard three-tier web application in AWS. The web and application tiers run on Amazon EC2 and the database tier runs on Amazon RDS. The company is redesigning the web and application tiers to use Amazon API Gateway and AWS Lambda, and the company intends to deploy the new application within 6 months. The IT Manager has asked the Solutions Architect to reduce costs in the interim.
Which solution will be MOST cost effective while maintaining reliability?

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

Answer: B

NEW QUESTION 116
An online retailer needs to regularly process large product catalogs, which are handled in batches. These are sent out to be processed by people using the Amazon Mechanical Turk service, but the retailer has asked its Solutions Architect to design a workflow orchestration system that allows it to handle multiple concurrent Mechanical Turk operations, deal with the result assessment process, and reprocess failures.
Which of the following options gives the retailer the ability to interrogate the state of every workflow with the LEAST amount of implementation effort?

A. Trigger Amazon CloudWatch alarms based upon message visibility in multiple Amazon SQS queues (one queue per workflow stage) and send messages via Amazon SNS to trigger AWS Lambda functions to process the next ste

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

Answer: C

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

NEW QUESTION 121
A financial services company is moving to AWS and wants to enable Developers to experiment and innovate while preventing access to production applications The company has the following requirements
• Production workloads cannot be directly connected to the internet
• All workloads must be restricted to the us-west-2 and eu-central-1 Regions
• Notification should be sent when Developer sandboxes exceed $500 in AWS spending monthly
Which combination of actions needs to be taken to create a multi-account structure that meets the company's requirements'? (Select THREE)

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

Answer: ABD

NEW QUESTION 124
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 127
A company is running a web application with On-Demand Amazon EC2 instances in Auto Scaling groups that scale dynamically based on custom metrics After extensive testing the company determines that the m5 2xlarge instance size is optimal for the workload Application data is stored in db r4 4xlarge Amazon RDS instances that are confirmed to be optimal The traffic to the web application spikes randomly during the day
What other cost-optimization methods should the company implement to further reduce costs without impacting the reliability of the application?

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

Answer: B

NEW QUESTION 131
A company is running a large application on-premises. Its technology stack consists of Microsoft .NET for the web server platform and Apache Cassandra for the

database. The company wants to migrate the application to AWS to improve service reliability. The IT team also wants to reduce the time it spends on capacity management and maintenance of this infrastructure. The Development team is willing and available to make code changes to support the migration.
Which design is the LEAST complex to manage after the migration?

A. Migrate the web servers to Amazon EC2 instances in an Auto Scaling group that is running .NE
B. Migrate the existing Cassandra database to Amazon Aurora with multiple read replicas, and run both in a Multi-AZ mode.
C. Migrate the web servers to an AWS Elastic Beanstalk environment that is running the .NET platform in a Multi-AZ Auto Scaling configuratio
D. Migrate the Cassandra database to Amazon EC2 instances that are running in a Multi-AZ configuration.
E. Migrate the web servers to an AWS Elastic Beanstalk environment that is running the .NET platform in a Multi-AZ Auto Scaling configuratio
F. Migrate the existing Cassandra database to Amazon DynamoDB.
G. Migrate the web servers to Amazon EC2 instances in an Auto Scaling group that is running .NE
H. Migrate the existing Cassandra database to Amazon DynamoDB.

Answer: B

NEW QUESTION 132
A company currently runs a secure application on Amazon EC2 that takes files from on-premises locations through AWS Direct Connect, processes them, and uploads them to a single Amazon S3 bucket. The application uses HTTPS for encryption in transit to Amazon S3, and S3 server-side encryption to encrypt at rest.
Which of the following changes should the Solutions Architect recommend to make this solution more secure without impeding application's performance?

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

Answer: D

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

NEW QUESTION 135
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 137
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 140
A company wants to replace its call system with a solution built using AWS managed services. The company call center would like the solution to receive calls, create contact flows, and scale to handle growth projections. The call center would also like the solution to use deep learning capabilities to recognize the intent of the callers and handle basic tasks, reducing the need to speak an agent. The solution should also be able to query business applications and provide relevant

information back to calls as requested.
Which services should the Solution Architect use to build this solution? (Choose three.)

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

Answer: BDE

NEW QUESTION 142
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 146
A company currently uses a single 1 Gbps AWS Direct Connect connection to establish connectivity between an AWS Region and its data center. The company has five Amazon VPCs, all of which are connected to the data center using the same Direct Connect connection. The Network team is worried about the single point of failure and is interested in improving the redundancy of the connections to AWS while keeping costs to a minimum.
Which solution would improve the redundancy of the connection to AWS while meeting the cost requirements?

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

Answer: B

NEW QUESTION 149
A Solutions Architect needs to design a highly available application that will allow authenticated users to stay connected to the application even when there are underlying failures
Which solution will meet these requirements?

A. Deploy the application on Amazon EC2 instances Use Amazon Route 53 to forward requests to the EC2 Instances Use Amazon DynamoDB to save the authenticated connection details
B. Deploy the application on Amazon EC2 instances in an Auto Scaling group Use an internet-facing Application Load Balancer to handle requests Use Amazon DynamoDB to save the authenticated connection details
C. Deploy the application on Amazon EC2 instances in an Auto Scaling group Use an internet-facing Application Load Balancer on the front end Use EC2 instances to save the authenticated connectiondetails
D. Deploy the application on Amazon EC2 instances in an Auto Scaling group Use an internet-facing Application Load Balancer on the front end Use EC2 instances hosting a MySQL database to save the authenticated connection details

Answer: B

NEW QUESTION 150
The Security team needs to provide a team of interns with an AWS environment so they can build the serverless video transcoding application. The project will use Amazon S3, AWS Lambda, Amazon API Gateway, Amazon Cognito, Amazon DynamoDB, and Amazon Elastic Transcoder.
The interns should be able to create and configure the necessary resources, but they may not have access to create or modify AWS IAM roles. The Solutions Architect creates a policy and attaches it to the interns' group.
How should the Security team configure the environment to ensure that the interns are self-sufficient?

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

H. The interns will provide the requirements for the permissions to be set in the role.

Answer: A

NEW QUESTION 153
A company is building an AWS landing zone and has asked a Solutions Architect to design a multi-account access strategy that will allow hundreds of users to use corporate credentials to access the AWS Console. The company is running a Microsoft Active Directory and users will use an AWS Direct Connect connection to connect to AWS. The company also wants to be able to federate to third-party services and providers, including custom applications.
Which solution meets the requirements by using the LEAST amount of management overhead?

A. Connect the Active Directory to AWS by using single sign-on and an Active Directory Federation Services (AD FS) with SAML 2.0, and then configure the identity Provider (IdP) system to use form-based authenticatio
B. Build the AD FS portal page with corporate branding, and integrate third-party applications that support SAML 2.0 as required.
C. Create a two-way Forest trust relationship between the on-premises Active Directory and the AWS Directory Servic
D. Set up AWS Single Sign-On with AWS Organization
E. Use single sign-on integrations for connections with third-party applications.
F. Configure single sign-on by connecting the on-premises Active Directory using the AWS Directory Service AD Connecto
G. Enable federation to the AWS services and accounts by using the IAM applications and services linking functio
H. Leverage third-party single sign-on as needed.
I. Connect the company's Active Directory to AWS by using AD FS and SAML 2.0. Configure the AD FS claim rule to leverage Regex and a common Active Directory naming convention for the security group to allow federation of all AWS account
J. Leverage third-party single sign-on as needed, and add it to the AD FS server.

Answer: D

Explanation:
https://aws.amazon.com/blogs/security/aws-federated-authentication-with-active-directory-federation-services-a

NEW QUESTION 156
A company wants to move a web application to AWS. The application stores session information locally on each web server, which will make auto scaling difficult. As part of the migration, the application will be rewritten to decouple the session data from the web servers. The company requires low latency, scalability, and availability.
Which service will meet the requirements for storing the session information in the MOST cost-effective way?

A. Amazon ElastiCache with the Memcached engine
B. Amazon S3
C. Amazon RDS MySQL
D. Amazon ElastiCache with the Redis engine

Answer: D

Explanation:
https://aws.amazon.com/caching/session-management/ https://aws.amazon.com/elasticache/redis-vs-memcached/

NEW QUESTION 159
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 162
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 166
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 167
A company has an application written using an in-house software framework. The framework installation takes 30 minutes and is performed with a user data script. Company Developers deploy changes to the application frequently. The framework installation is becoming a bottleneck in this process.
Which of the following would speed up this process?

A. Create a pipeline to build a custom AMI with the framework installed and use this AMI as a baseline for application deployments.
B. Employ a user data script to install the framework but compress the installation files to make them smaller.
C. Create a pipeline to parallelize the installation tasks and call this pipeline from a user data script.
D. Configure an AWS OpsWorks cookbook that installs the framework instead of employing user dat
E. Use this cookbook as a base for all deployments.

Answer: A

Explanation:
https://aws.amazon.com/codepipeline/features/?nc=sn&loc=2

NEW QUESTION 172
A company is running multiple applications on Amazon EC2. Each application is deployed and managed by multiple business units. All applications are deployed on a single AWS account but on different virtual private clouds (VPCs). The company uses a separate VPC in the same account for test and development purposes.
Production applications suffered multiple outages when users accidentally terminated and modified resources that belonged to another business unit. A Solutions Architect has been asked to improve the availability of the company applications while allowing the Developers access to the resources they need.
Which option meets the requirements with the LEAST disruption?

A. Create an AWS account for each business uni
B. Move each business unit's instances to its own account and set up a federation to allow users to access their business unit's account.
C. Set up a federation to allow users to use their corporate credentials, and lock the users down to their own VP
D. Use a network ACL to block each VPC from accessing other VPCs.
E. Implement a tagging policy based on business unit
F. Create an IAM policy so that each user can terminate instances belonging to their own business units only.
G. Set up role-based access for each user and provide limited permissions based on individual roles and the services for which each user is responsible.

Answer: C

Explanation:
Principal – Control what the person making the request (the principal) is allowed to do based on the tags that are attached to that person's IAM user or role. To do this, use the aws:PrincipalTag/key-name condition key to specify what tags must be attached to the IAM user or role before the request is allowed.
https://docs.aws.amazon.com/IAM/latest/UserGuide/access_iam-tags.html

NEW QUESTION 176
A company collects a steady stream of 10 million data records from 100,000 sources each day. These records are written to an Amazon RDS MySQL DB. A query must produce the daily average of a data source over the past 30 days. There are twice as many reads as writes. Queries to the collected data are for one source ID at a time.
How can the Solutions Architect improve the reliability and cost effectiveness of this solution?

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

Answer: B

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

NEW QUESTION 180
A company has developed a web application that runs on Amazon EC2 instances in one AWS Region. The company has taken on new business in other countries and must deploy its application into other to meet low-latency requirements for its users. The regions can be segregated, and an application running in one region does not need to communicate with instances in other regions.
How should the company's Solutions Architect automate the deployment of the application so that it can be MOST efficiently deployed into multiple regions?

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

Answer: D

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

NEW QUESTION 181
A bank is re-architecting its mainframe-based credit card approval processing application to a cloud-native application on the AWS cloud.
The new application will receive up to 1,000 requests per second at peak load. There are multiple steps to each transaction, and each step must receive the result of the previous step. The entire request must return an authorization response within less than 2 seconds with zero data loss. Every request must receive a response. The solution must be Payment Card Industry Data Security Standard (PCI DSS)-compliant.
Which option will meet all of the bank's objectives with the LEAST complexity and LOWEST cost while also meeting compliance requirements?

A. Create an Amazon API Gateway to process inbound requests using a single AWS Lambda task that performs multiple steps and returns a JSON object with the approval statu
B. Open a support case to increase the limit for the number of concurrent Lambdas to allow room for bursts of activity due to the new application.
C. Create an Application Load Balancer with an Amazon ECS cluster on Amazon EC2 Dedicated instances in a target group to process incoming request
D. Use Auto Scaling to scale the cluster out/in based on average CPU utilizatio
E. Deploy a web service that processes all of the approval steps and returns a JSON object with the approval status.
F. Deploy the application on Amazon EC2 on Dedicated Instance
G. Use an Elastic Load Balancer in front of a farm of application servers in an Auto Scaling group to handle incoming request
H. Scale out/in based on a custom Amazon CloudWatch metric for the number of inbound requests per second after measuring the capacity of a single instance.
I. Create an Amazon API Gateway to process inbound requests using a series of AWS Lambda processes, each with an Amazon SQS input queu
J. As each step completes, it writes its result to the next step's queu
K. The final step returns a JSON object with the approval statu
L. Open a support case to increase the limit for the number of concurrent Lambdas to allow room for bursts of activity due to the new application.

Answer: B

NEW QUESTION 183
A Company has a security event whereby an Amazon S3 bucket with sensitive information was made public. Company policy is to never have public S3 objects, and the Compliance team must be informed immediately when any public objects are identified.
How can the presence of a public S3 object be detected, set to trigger alarm notifications, and automatically remediated in the future? (Choose two.)

A. Turn on object-level logging for Amazon S3. Turn on Amazon S3 event notifications to notify by using an Amazon SNS topic when a PutObject API call is made with a public-read permission.
B. Configure an Amazon CloudWatch Events rule that invokes an AWS Lambda function to secure the S3 bucket.
C. Use the S3 bucket permissions for AWS Trusted Advisor and configure a CloudWatch event to notify by using Amazon SNS.
D. Turn on object-level logging for Amazon S3. Configure a CloudWatch event to notify by using an SNS topic when a PutObject API call with public-read permission is detected in the AWS CloudTrail logs.
E. Schedule a recursive Lambda function to regularly change all object permissions inside the S3 bucket.

Answer: BD

Explanation:
https://aws.amazon.com/blogs/security/how-to-detect-and-automatically-remediate-unintended-permissions-in-a

NEW QUESTION 188
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 190
......

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