Amazon-Web-Services

Exam Questions SAP-C01

AWS Certified Solutions Architect- Professional

NEW QUESTION 1
A group of Amazon EC2 instances have been configured as high performance computing (HPC) cluster. The instances are running in a placement group, and are able to communicate with each other at network of up to 20 Gbps.
The cluster needs to communicate with a control EC2 instance outside of the placement group. The control instance has the same instance type and AMI as the other instances, and is configured with a public IP address.
How can the Solutions Architect improve the network speeds between the control instance and the instances in the placement group?

A. Terminate the control instance and relaunch in the placement group.
B. Ensure that the instances are communicating using the private IP addresses.
C. Ensure that the control instance is using an Elastic Network Adapter.
D. Move the control instance inside the placement group.

Answer: D

Explanation:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html

NEW QUESTION 2
A company receives clickstream data files to Amazon S3 every five minutes. A Python script runs as a cron job once a day on an Amazon EC2 instance to process each file and load it into a database hosted on Amazon RDS. The cron job takes 15 to 30 minutes to process 24 hours of data. The data consumers ask for the data be available as soon as possible.
Which solution would accomplish the desired outcome?

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

Answer: D

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

NEW QUESTION 3
A retail company processes point-of-state data on application servers in its data center and writes outputs to Amazon DynamoDB table. The data center is connected to the company's VPC with an AWS Direct Connect (DX) connection, and the application servers require a consistent network connection at speed greater than 2 Gbps.
The company decides that the DynamoDB table needs to be highly available and fault tolerant. The company policy states that the data should be available across two regions.
What changes should the company make to meet these requirements?

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

Answer: A

NEW QUESTION 4
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 5
A company runs a memory-intensive analytics application using on-demand Amazon EC2 compute optimized instance. The application is used continuously and application demand doubles during working hours. The application currently scales based on CPU usage. When scaling in occurs, a lifecycle hook is used because the instance requires 4 minutes to clean the application state before terminating.
Because users reported poor performance during working hours, scheduled scaling actions were implemented so additional instances would be added during working hours. The Solutions Architect has been asked to reduce the cost of the application.
Which solution is MOST cost-effective?

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

Answer: D

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

NEW QUESTION 6
An on-premises application will be migrated to the cloud. The application consists of a single Elasticsearch virtual machine with data source feeds from local systems that will not be migrated, and a Java web application on Apache Tomcat running on three virtual machines. The Elasticsearch server currently uses 1 TB of storage out of 16 TB available storage, and the web application is updated every 4 months. Multiple users access the web application from the Internet. There is a 10Gbit AWS Direct Connect connection established, and the application can be migrated over a schedules 48-hour change window.
Which strategy will have the LEAST impact on the Operations staff after the migration?

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

Answer: D

NEW QUESTION 7
A company uses an Amazon EMR cluster to process data once a day. The raw data comes from Amazon S3, and the resulting processed data is also stored in Amazon S3. The processing must complete within 4 hours; currently, it only takes 3 hours. However, the processing time is taking 5 to 10 minutes. longer each week due to an increasing volume of raw data.
The team is also concerned about rising costs as the compute capacity increases. The EMR cluster is currently running on three m3.xlarge instances (one master and two core nodes).
Which of the following solutions will reduce costs related to the increasing compute needs?

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

Answer: B

NEW QUESTION 8
While debugging a backend application for an IoT system that supports globally distributed devices a Solutions Architect notices that stale data is occasionally being sent to user devices. Devices often share data, and stale data does not cause issues in most cases However device operations are disrupted when a device reads the stale data after an update
The global system has multiple identical application stacks deployed In different AWS Regions If a user device travels out of its home geographic region it will always connect to the geographically closest AWS Region to write or read data The same data is available in all supported AWS Regions using an Amazon DynamoDB global table
What change should be made to avoid causing disruptions in device operations'?

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

Answer: B

NEW QUESTION 9
A company has a requirement that only allows specially hardened AMIs to be launched into public subnets in a VPC, and for the AMIs to be associated with a specific security group. Allowing non-compliant instances to launch into the public subnet could present a significant security risk if they are allowed to operate. A mapping of approved AMIs to subnets to security groups exists in an Amazon DynamoDB table in the same AWS account. The company created an AWS Lambda function that, when invoked, will terminate a given Amazon EC2 instance if the combination of AMI, subnet, and security group are not approved in the DynamoDB table.
What should the Solutions Architect do to MOST quickly mitigate the risk of compliance deviations?

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

Answer: C

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

NEW QUESTION 10
A company has a serverless application comprised of Amazon CloudFront, Amazon API Gateway, and AWS Lambda functions. The current deployment process of the application code is to create a new version number of the Lambda function and run an AWS CLI script to update. If the new function version has errors, another CLI script reverts by deploying the previous working version of the function. The company would like to decrease the time to deploy new versions of the application logic provided by the Lambda functions, and also reduce the time to detect and revert when errors are identified.
How can this be accomplished?

A. Create and deploy nested AWS CloudFormation stacks with the parent stack consisting of the AWS CloudFront distribution and API Gateway, and the child stack containing the Lambda functio
B. For changes to Lambda, create an AWS CloudFormation change set and deploy; if errors are triggered, revert the AWS CloudFormation change set to the previous version.
C. Use AWS SAM and built-in AWS CodeDeploy to deploy the new Lambda version, gradually shift traffic to the new version, and use pre-traffic and post-traffic test functions to verify cod
D. Rollback if Amazon CloudWatch alarms are triggered.
E. Refactor the AWS CLI scripts into a single script that deploys the new Lambda versio
F. When deployment is completed, the script tests execut
G. If errors are detected, revert to the previous Lambda version.
H. Create and deploy an AWS CloudFormation stack that consists of a new API Gateway endpoint that references the new Lambda versio
I. Change the CloudFront origin to the new API Gateway endpoint, monitor errors and if detected, change the AWS CloudFront origin to the previous API Gateway endpoint.

Answer: B

Explanation:
https://aws.amazon.com/about-aws/whats-new/2017/11/aws-lambda-supports-traffic-shifting-and-phased-deploy https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/automating-updates-to-serverle

NEW QUESTION 10
A company would like to implement a serverless application by using Amazon API Gateway, AWS Lambda and Amazon DynamoDB. They deployed a proof of concept and stated that the average response time is greater than what their upstream services can accept Amazon CloudWatch metrics did not indicate any issues with DynamoDB but showed that some Lambda functions were hitting their timeout.
Which of the following actions should the Solutions Architect consider to improve performance? (Choose two.)

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

Answer: BD

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

NEW QUESTION 15
A large company has increased its utilization of AWS over time in an unmanaged way. As such, they have a large number of independent AWS accounts across different business units, projects, and environments. The company has created a Cloud Center of Excellence team, which is responsible for managing all aspects of the AWS Cloud, including their AWS accounts.
Which of the following should the Cloud Center of Excellence team do to BEST address their requirements in a centralized way? (Select two.)

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

Answer: DE

NEW QUESTION 18
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 23
A company is moving a business-critical, multi-tier application to AWS. The architecture consists of a desktop client application and server infrastructure. The server infrastructure resides in an on-premises data center that frequently fails to maintain the application uptime SLA of 99.95%. A Solutions Architect must re-architect the application to ensure that it can meet or exceed the SLA.
The application contains a PostgreSQL database running on a single virtual machine. The business logic and presentation layers are load balanced between multiple virtual machines. Remote users complain about slow load times while using this latency-sensitive application.
Which of the following will meet the availability requirements with little change to the application while improving user experience and minimizing costs?

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

Answer: B

NEW QUESTION 27
A company is running an email application across multiple AWS Regions. The company uses Ohio (us-east-2) as the primary Region and Northern Virginia (us-east-1) as the Disaster Recovery (DR) Region. The data is continuously replicated from the primary Region to the DR Region by a single instance on the public subnet in both Regions. The replication messages between the Regions have a significant backlog during certain times of the day. The backlog clears on its own after a short time, but it affects the application's RPO.
Which of the following solutions should help remediate this performance problem? (Select TWO)

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

Answer: AC

NEW QUESTION 32
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 35
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 37
A company has an Amazon VPC that is divided into a public subnet and a private subnet A web application runs in Amazon VPC, and each subnet has its own NACL The public subnet has a CIDR of 10.0.0.0/24. An Application Load Balancer is deployed to the public subnet. The private subnet has a CIDR of 10.0.1.0/24. Amazon EC2 instances that run a web server on port 80 are launched into the private subnet.
Only network traffic that is required for the Application Load Balancer to access the web application can be allowed to travel between the public and private subnets
What collection of rules should be written to ensure that the private subnet's NACL meets the requirement? (Select TWO.)

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

Answer: BE

NEW QUESTION 39
A large multinational company runs a timesheet application on AWS that is used by staff across the world. The application runs on Amazon EC2 instances in an Auto Scaling group behind an Elastic Load Balancing (ELB) load balancer, and stores in an Amazon RDS MySQL Multi-AZ database instance.
The CFO is concerned about the impact on the business if the application is not available. The application must not be down for more than two hours, but the solution must be as cost-effective as possible.
How should the Solutions Architect meet the CFO's requirements while minimizing data loss?

A. In another region, configure a read replica and create a copy of the infrastructur
B. When an issue occurs, promote the read replica and configure as an Amazon RDS Multi-AZ database instanc
C. Update the DNS to point to the other region's ELB.
D. Configure a 1-day window of 60-minute snapshots of the Amazon RDS Multi-AZ database instance.Create an AWS CloudFormation template of the application

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

Answer: D

NEW QUESTION 40
A company is migrating an application to AWS. It wants to use fully managed services as much as possible during the migration. The company needs to store large, important documents within the application with the following requirements:

> The data must be highly durable and available.

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

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

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

Answer: B

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

NEW QUESTION 42
A company is using AWS for production and development workloads. Each business unit has its own AWS account for production, and a separate AWS account to develop and deploy its applications. The Information Security department has introduced new security policies that limit access for terminating certain Amazon ECs instances in all accounts to a small group of individuals from the Security team.
How can the Solutions Architect meet these requirements?

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

Answer: B

NEW QUESTION 45
A company runs an ordering system on AWS using Amazon SQS and AWS Lambda, with each order received as a JSON message. recently the company had a marketing event that led to a tenfold increase in orders. With this increase, the following undesired behaviors started in the ordering system:

> Lambda failures while processing orders lead to queue backlogs.

> The same orders have been processed multiple times.
A solutions Architect has been asked to solve the existing issues with the ordering system and add the following resiliency features:

> Retain problematic orders for analysis.

> Send notification if errors go beyond a threshold value. How should the Solutions Architect meet these requirements?

A. Receive multiple messages with each Lambda invocation, add error handling to message processing code and delete messages after processing, increase the visibility timeout for the messages, create a dead letter queue for messages that could not be processed, create an Amazon CloudWatch alarm on Lambda errors for notification.
B. Receive single messages with each Lambda invocation, put additional Lambda workers to poll the queue, delete messages after processing, increase the message timer for the messages, use Amazon CloudWatch Logs for messages that could not be processed, create a CloudWatch alarm on Lambda errors for notification.
C. Receive multiple messages with each Lambda invocation, use long polling when receiving the messages, log the errors from the message processing code using Amazon CloudWatch Logs, create a dead letter queue with AWS Lambda to capture failed invocations, create CloudWatch events on Lambda errors for notification.
D. Receive multiple messages with each Lambda invocation, add error handling to message processing code and delete messages after processing, increase the visibility timeout for the messages, create a delay queue for messages that could not be processed, create an Amazon CloudWatch metric on Lambda errors for notification.

Answer: D

NEW QUESTION 46
A company runs an IoT platform on AWS. IoT sensors in various locations send data to the company's Node.js API servers on Amazon EC2 instances running behind an Application Load Balancer. The data is stored in an Amazon RDS MySQL DB instance that uses a 4 TB General Purpose SSD volume.
The number of sensors the company has deployed in the field has increased over time, and is expected to grow significantly. The API servers are consistently overloaded and RDS metrics show high write latency.
Which of the following steps together will resolve the issues permanently and enable growth as new sensors are provisioned, while keeping this platform cost-efficient? (Choose two.)

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

Answer: CE

NEW QUESTION 49
A large company experienced a drastic increase in its monthly AWS spend. This is after Developers accidentally launched Amazon EC2 instances in unexpected regions. The company has established practices around least privileges for Developers and controls access to on-premises resources using Active Directory groups. The company now wants to control costs by restricting the level of access that Developers have to the AWS Management Console without impacting their productivity. The company would also like to allow Developers to launch Amazon EC2 in only one region, without limiting access to other services in any region.
How can this company achieve these new security requirements while minimizing the administrative burden on the Operations team?

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

Answer: D

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

NEW QUESTION 52
A company runs a Windows Server host in a public subnet that is configured to allow a team of administrators to connect over RDP to troubleshoot issues with hosts in a private subnet. The host must be available at all times outside of a scheduled maintenance window, and needs to receive the latest operating system updates within 3 days of release.
What should be done to manage the host with the LEAST amount of administrative effort?

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

Answer: B

NEW QUESTION 53
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 54
AnyCompany has acquired numerous companies over the past few years. The CIO for AnyCompany would like to keep the resources for each acquired company separate. The CIO also would like to enforce a chargeback model where each company pays for the AWS services it uses.
The Solutions Architect is tasked with designing an AWS architecture that allows AnyCompany to achieve the following:

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

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

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

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

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

A. Create a multi-account strategy with an account per compan
B. Use consolidated billing to ensure that AnyCompany needs to pay a single bill only.
C. Create a multi-account strategy with a virtual private cloud (VPC) for each compan
D. Reduce impact across companies by not creating any VPC peering link
E. As everything is in a single account, there will be a single invoic
F. use tagging to create a detailed bill for each company.
G. Create IAM users for each Developer in the account to which they require acces
H. Create policies that allow the users access to all resources in that accoun
I. Attach the policies to the IAM user.
J. Create a federated identity store against the company's Active Director
K. Create IAM roles with appropriate permissions and set the trust relationships with AWS and the identity stor
L. Use AWS STS to grant users access based on the groups they belong to in the identity store.
M. Create a multi-account strategy with an account per compan
N. For billing purposes, use a tagging solution that uses a tag to identify the company that creates each resource.

Answer: AD

NEW QUESTION 57
A company is migrating to the cloud. It wants to evaluate the configurations of virtual machines in its existing data center environment to ensure that it can size new Amazon EC2 instances accurately. The company wants to collect metrics, such as CPU, memory, and disk utilization, and it needs an inventory of what processes are running on each instance. The company would also like to monitor network connections to map communications between servers.
Which would enable the collection of this data MOST cost effectively?

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

Answer: A

NEW QUESTION 60
A Solutions Architect must create a cost-effective backup solution for a company's 500MB source code repository of proprietary and sensitive applications. The repository runs on Linux and backs up daily to tape. Tape backups are stored for 1 year.
The current solutions are not meeting the company's needs because it is a manual process that is prone to error, expensive to maintain, and does not meet the need for a Recovery Point Objective (RPO) of 1 hour or Recovery Time Objective (RTO) of 2 hours. The new disaster recovery requirement is for backups to be stored offsite and to be able to restore a single file if needed.
Which solution meets the customer's needs for RTO, RPO, and disaster recovery with the LEAST effort and expense?

A. Replace local tapes with an AWS Storage Gateway virtual tape library to integrate with current backup softwar
B. Run backups nightly and store the virtual tapes on Amazon S3 standard storage inUS-EAST-1. Use cross-region replication to create a second copy in US-WEST-2. Use Amazon S3 lifecycle policies to perform automatic migration to Amazon Glacier and deletion of expired backups after 1 year?
C. Configure the local source code repository to synchronize files to an AWS Storage Gateway file Amazon gateway to store backup copies in an Amazon S3 Standard bucke
D. Enable versioning on the Amazon S3 bucke
E. Create Amazon S3 lifecycle policies to automatically migrate old versions of objects to Amazon S3 Standard 0 Infrequent Access, then Amazon Glacier, then delete backups after 1 year.
F. Replace the local source code repository storage with a Storage Gateway stored volum
G. Change the default snapshot frequency to 1 hou
H. Use Amazon S3 lifecycle policies to archive snapshots to Amazon Glacier and remove old snapshots after 1 yea
I. Use cross-region replication to create a copy of the snapshots in US-WEST-2.
J. Replace the local source code repository storage with a Storage Gateway cached volum
K. Create a snapshot schedule to take hourly snapshot
L. Use an Amazon CloudWatch Events schedule expression rule to run on hourly AWS Lambda task to copy snapshots from US-EAST -1 to US-WEST-2.

Answer: B

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

NEW QUESTION 61
A company's data center is connected to the AWS Cloud over a minimally used 10-Gbps AWS Direct Connect connection with a private virtual interface to its virtual private cloud {VPC) The company internet connection is 200 Mbps and the company has a 150-TB dataset that is created each Friday The data must be transferred and available in Amazon S3 on Monday morning
Which is the LEAST expensive way to meet the requirements while allowing for data transfer growth?

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

Answer: D

NEW QUESTION 65
A company has a large on-premises Apache Hadoop cluster with a 20 PB HDFS database. The cluster is growing every quarter by roughly 200 instances and 1 PB. The company's goals are to enable resiliency for its Hadoop data, limit the impact of losing cluster nodes, and significantly reduce costs. The current cluster runs 24/7 and supports a variety of analysis workloads, including interactive queries and batch processing.
Which solution would meet these requirements with the LEAST expense and down time?

A. Use AWS Snowmobile to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster initially sized to handle the interactive workload based on historical data from theon-premises cluste
B. Store the data on EMRF
C. Minimize costs using Reserved Instances for master and core nodes and Spot Instances for task nodes, and auto scale task nodes based on Amazon CloudWatch metric
D. Create job-specific, optimized clusters for batch workloads that are similarly optimized.
E. Use AWS Snowmobile to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster of similar size and configuration to the current cluste
F. Store the data on EMRF
G. Minimize costs by using Reserved Instance
H. As the workload grows each quarter, purchase additional Reserved Instances and add to the cluster.
I. Use AWS Snowball to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster initially sized to handle the interactive workloads based on historical data from theon-premises cluste
J. Store the on EMRF
K. Minimize costs using Reserved Instances for master and core nodes and Spot Instances for task nodes, and auto scale task nodes based on Amazon CloudWatch metric
L. Create job-specific, optimized clusters for batch workloads that are similarly optimized.
M. Use AWS Direct Connect to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster initially sized to handle the interactive workload based on historical data from theon-premises cluste
N. Store the data on EMRF
O. Minimize costs using Reserved Instances for master and core nodes and Spot Instances for task nodes, and auto scale task nodes based on Amazon CloudWatch metric
P. Create job-specific, optimized clusters for batch workloads that are similarly optimized.

Answer: A

Explanation:
Q: How should I choose between Snowmobile and Snowball?
To migrate large datasets of 10PB or more in a single location, you should use Snowmobile. For datasets less than 10PB or distributed in multiple locations, you should use Snowball. In addition, you should evaluate the amount of available bandwidth in your network backbone. If you have a high speed backbone with hundreds of Gb/s of spare throughput, then you can use Snowmobile to migrate the large datasets all at once. If you have limited bandwidth on your backbone, you should consider using multiple Snowballs to migrate the data incrementally.

NEW QUESTION 68
A Solutions Architect must migrate an existing on-premises web application with 70 TB of static files supporting a public open-data initiative. The architect wants to upgrade to the latest version of the host operating system as part of the migration effort.
Which is the FASTEST and MOST cost-effective way to perform the migration?

A. Run a physical-to-virtual conversion on the application serve
B. Transfer the server image over the internet, and transfer the static data to Amazon S3.
C. Run a physical-to-virtual conversion on the application serve
D. Transfer the server image over AWS Direct Connect, and transfer the static data to Amazon S3.
E. Re-platform the server to Amazon EC2, and use AWS Snowball to transfer the static data to Amazon S3.
F. Re-platform the server by using the AWS Server Migration Service to move the code and data to a new Amazon EC2 instance.

Answer: C

NEW QUESTION 73
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 76
An organization has a write-intensive mobile application that uses Amazon API Gateway, AWS Lambda, and Amazon DynamoDB. The application has scaled well, however, costs have increased exponentially because of higher than anticipated Lambda costs. The application's use is unpredictable, but there has been a steady 20% increase in utilization every month.
While monitoring the current Lambda functions, the Solutions Architect notices that the execution-time averages 4.5 minutes. Most of the wait time is the result of a high-latency network call to a 3-TB MySQL database server that is on-premises. A VPN is used to connect to the VPC, so the Lambda functions have been configured with a five-minute timeout.
How can the Solutions Architect reduce the cost of the current architecture?

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

Answer: D

NEW QUESTION 77
A Solution Architect is designing a deployment strategy for an application tier and gas the following requirements.
* The application code will need a 500 HB static dataset to be present before application startup.
* The application tier be able to scale Up and down based on demand with as little startup time as possible.
* The development team should be able to update the code multiple times each day.
* Critical operating system (OS) patches must be installed within 48 hours of being released. Which deployment strategy meets these requirements?

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

Answer: C

NEW QUESTION 79
To abide by industry regulations, a Solutions Architect must design a solution that will store a company's critical data in multiple public AWS Regions, including in the United States, where the company's headquarters is located. The Solutions Architect is required to provide access to the data stored in AWS to the company's global WAN network. The Security team mandates that no traffic accessing this data should traverse the public internet.
How should the Solutions Architect design a highly available solution that meets the requirements and is cost-effective?

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

Answer: D

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

NEW QUESTION 84
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 89
A retail company is running an application that stores invoice files in Amazon S3 bucket and metadata about the files in an Amazon. The S3 bucket and DynamoDB table are in us-east-1. The company wants to protect itself from data corruption and loss of connectivity to either Region.
Which option meets these requirements?

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

Answer: A

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

NEW QUESTION 93
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 94
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 98
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 103
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 107
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 112
A company has asked a Solutions Architect to design a secure content management solution that can be accessed by API calls by external customer applications. The company requires that a customer administrator must be able to submit an API call and roll back changes to existing files sent to the content management solution, as needed.
What is the MOST secure deployment design that meets all solution requirements?

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

Answer: A

NEW QUESTION 116
A Solutions Architect needs to design a highly available application that will allow authenticated users to stay connected to the application even when there are underlying failures
Which solution will meet these requirements?

A. Deploy the application on Amazon EC2 instances Use Amazon Route 53 to forward requests to the EC2 Instances Use Amazon DynamoDB to save the authenticated connection details
B. Deploy the application on Amazon EC2 instances in an Auto Scaling group Use an internet-facing Application Load Balancer to handle requests Use Amazon DynamoDB to save the authenticated connection details
C. Deploy the application on Amazon EC2 instances in an Auto Scaling group Use an internet-facing Application Load Balancer on the front end Use EC2 instances to save the authenticated 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 117
A large company is migrating its entire IT portfolio to AWS. Each business unit in the company has a standalone AWS account that supports both development and test environments. New accounts to support production workloads will be needed soon.
The Finance department requires a centralized method for payment but must maintain visibility into each group's spending to allocate costs.
The Security team requires a centralized mechanism to control IAM usage in all the company's accounts. What combination of the following options meet the

company's needs with LEAST effort? (Choose two.)

A. Use a collection of parameterized AWS CloudFormation templates defining common IAM permissions that are launched into each accoun
B. Require all new and existing accounts to launch the appropriate stacks to enforce the least privilege model.
C. Use AWS Organizations to create a new organization from a chosen payer account and define an organizational unit hierarch
D. Invite the existing accounts to join the organization and create new accounts using Organizations.
E. Require each business unit to use its own AWS account
F. Tag each AWS account appropriately and enable Cost Explorer to administer chargebacks.
G. Enable all features of AWS Organizations and establish appropriate service control policies that filter IAM permissions for sub-accounts.
H. Consolidate all of the company's AWS accounts into a single AWS accoun
I. Use tags for billing purposes and IAM's Access Advice feature to enforce the least privilege model.

Answer: BD

NEW QUESTION 120
A company is planning the migration of several lab environments used for software testing. An assortment of custom tooling is used to manage the test runs for each lab. The labs use immutable infrastructure for the software test runs, and the results are stored in a highly available SQL database cluster. Although completely rewriting the custom tooling is out of scope for the migration project, the company would like to optimize workloads during the migration.
Which application migration strategy meets this requirement?

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

Answer: B

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

NEW QUESTION 124
A company has deployed an application to multiple environments in AWS, including production and testing. The company has separate accounts for production and testing, and users are allowed to create additional application users for team members or services, as needed. The Security team has asked the Operations team for better isolation between production and testing with centralized controls on security credentials and improved management of permissions between environments.
Which of the following options would MOST securely accomplish this goal?

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

Answer: A

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

NEW QUESTION 126
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 128

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

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

Answer: D

Explanation:
https://aws.amazon.com/s3/faqs/

NEW QUESTION 129
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 131
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 132
......

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

SAP-C01 Practice Exam Features:

* SAP-C01 Questions and Answers Updated Frequently

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

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

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

100% Actual & Verified — Instant Download, Please Click
Order The SAP-C01 Practice Test Here