

SAP-C02 Dumps

AWS Certified Solutions Architect - Professional

<https://www.certleader.com/SAP-C02-dumps.html>



NEW QUESTION 1

- (Exam Topic 1)

A startup company recently migrated a large ecommerce website to AWS. The website has experienced a 70% increase in sales. Software engineers are using a private GitHub repository to manage code. The DevOps learn is using Jenkins for builds and unit testing. The engineers need to receive notifications for bad builds and zero downtime during deployments. The engineers also need to ensure any changes to production are seamless for users and can be rolled back in the event of a major issue.

The software engineers have decided to use AWS CodePipeline to manage their build and deployment process.

Which solution will meet these requirements?

- A. Use GitHub websockets to trigger the CodePipeline pipeline
- B. Use the Jenkins plugin for AWS CodeBuild to conduct unit testing
- C. Send alerts to an Amazon SNS topic for any bad build
- D. Deploy in an in-place
- E. all-at-once deployment configuration using AWS CodeDeploy.
- F. Use GitHub webhooks to trigger the CodePipeline pipeline
- G. Use the Jenkins plugin for AWS CodeBuild to conduct unit testing
- H. Send alerts to an Amazon SNS topic for any bad build
- I. Deploy in a blue/green deployment using AWS CodeDeploy.
- J. Use GitHub websockets to trigger the CodePipeline pipeline
- K. Use AWS X-Ray for unit testing and static code analysis
- L. Send alerts to an Amazon SNS topic for any bad build
- M. Deploy in a blue/green deployment using AWS CodeDeploy.
- N. Use GitHub webhooks to trigger the CodePipeline pipeline
- O. Use AWS X-Ray for unit testing and static code analysis
- P. Send alerts to an Amazon SNS topic for any bad build
- Q. Deploy in an in-place, all-at-once deployment configuration using AWS CodeDeploy.

Answer: B

NEW QUESTION 2

- (Exam Topic 1)

A solutions architect is designing a publicly accessible web application that is on an Amazon CloudFront distribution with an Amazon S3 website endpoint as the origin. When the solution is deployed, the website returns an Error 403: Access Denied message.

Which steps should the solutions architect take to correct the issue? (Select TWO.)

- A. Remove the S3 block public access option from the S3 bucket.
- B. Remove the requester pays option from the S3 bucket.
- C. Remove the origin access identity (OAI) from the CloudFront distribution.
- D. Change the storage class from S3 Standard to S3 One Zone-Infrequent Access (S3 One Zone-IA).
- E. Disable S3 object versioning.

Answer: AB

Explanation:

See using S3 to host a static website with Cloudfront: <https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-serve-static-website/>

- Using a REST API endpoint as the origin, with access restricted by an origin access identity (OAI)
- Using a website endpoint as the origin, with anonymous (public) access allowed
- Using a website endpoint as the origin, with access restricted by a Referer header

NEW QUESTION 3

- (Exam Topic 1)

A company wants to deploy an AWS WAF solution to manage AWS WAF rules across multiple AWS accounts. The accounts are managed under different OUs in AWS Organizations.

Administrators must be able to add or remove accounts or OUs from managed AWS WAF rule sets as needed. Administrators also must have the ability to automatically update and remediate noncompliant AWS WAF rules in all accounts

Which solution meets these requirements with the LEAST amount of operational overhead?

- A. Use AWS Firewall Manager to manage AWS WAF rules across accounts in the organization
- B. Use an AWS Systems Manager Parameter Store parameter to store account numbers and OUs to manage Update the parameter as needed to add or remove accounts or OUs Use an Amazon EventBridge (Amazon CloudWatch Events) rule to identify any changes to the parameter and to invoke an AWS Lambda function to update the security policy in the Firewall Manager administrative account
- C. Deploy an organization-wide AWS Config rule that requires all resources in the selected OUs to associate the AWS WAF rule
- D. Deploy automated remediation actions by using AWS Lambda to fix noncompliant resource
- E. Deploy AWS WAF rules by using an AWS CloudFormation stack set to target the same OUs where the AWS Config rule is applied.
- F. Create AWS WAF rules in the management account of the organization
- G. Use AWS Lambda environment variables to store account numbers and OUs to manage Update environment variables as needed to add or remove accounts or OUs Create cross-account IAM roles in member account
- H. Assume the roles by using AWS Security Token Service (AWS STS) in the Lambda function to create and update AWS WAF rules in the member accounts
- I. Use AWS Control Tower to manage AWS WAF rules across accounts in the organization
- J. Use AWS Key Management Service (AWS KMS) to store account numbers and OUs to manage Update AWS KMS as needed to add or remove accounts or OUs
- K. Create IAM users in member accounts Allow AWS Control Tower in the management account to use the access key and secret access key to create and update AWS WAF rules in the member accounts

Answer: B

NEW QUESTION 4

- (Exam Topic 1)

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 instance
- B. Determine the minimum number of website instances required during off-peak times and use Spot Instances to cover them while using Reserved Instances to cover peak demand
- C. Use Amazon EC2 R4 and Amazon EC2 R5 Reserved Instances in an Auto Scaling group for the video analysis application
- D. Keep the website on T2 instance
- E. 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
- F. Use Spot Fleet for the video analysis application comprised of Amazon EC2 C4 and Amazon EC2 C5 Spot Instances.
- G. Migrate the website to AWS Elastic Beanstalk and Amazon EC2 C4 instance
- H. 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.
- I. Migrate the website to AWS Elastic Beanstalk and Amazon EC2 R4 instance
- J. 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 5

- (Exam Topic 1)

A company hosts a large on-premises MySQL database at its main office that supports an issue tracking system used by employees around the world. The company already uses AWS for some workloads and has created an Amazon Route 53 entry for the database endpoint that points to the on-premises database. Management is concerned about the database being a single point of failure and wants a solutions architect to migrate the database to AWS without any data loss or downtime.

Which set of actions should the solutions architect implement?

- A. Create an Amazon Aurora DB cluster
- B. Use AWS Database Migration Service (AWS DMS) to do a full load from the on-premises database to Aurora
- C. Update the Route 53 entry for the database to point to the Aurora cluster endpoint
- D. and shut down the on-premises database.
- E. During nonbusiness hours, shut down the on-premises database and create a backup
- F. Restore this backup to an Amazon Aurora DB cluster
- G. When the restoration is complete, update the Route 53 entry for the database to point to the Aurora cluster endpoint, and shut down the on-premises database.
- H. Create an Amazon Aurora DB cluster
- I. Use AWS Database Migration Service (AWS DMS) to do a full load with continuous replication from the on-premises database to Aurora
- J. When the migration is complete, update the Route 53 entry for the database to point to the Aurora cluster endpoint, and shut down the on-premises database.
- K. Create a backup of the database and restore it to an Amazon Aurora multi-master cluster
- L. This Aurora cluster will be in a master-master replication configuration with the on-premises database
- M. Update the Route 53 entry for the database to point to the Aurora cluster endpoint
- N. and shut down the on-premises database.

Answer: C

Explanation:

"Around the world" eliminates possibility for the maintenance window at night. The other difference is ability to leverage continuous replication in MySQL to Aurora case.

NEW QUESTION 6

- (Exam Topic 1)

A company is running an application distributed over several Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer. The security team requires that all application access attempts be made available for analysis. Information about the client IP address, connection type, and user agent must be included.

Which solution will meet these requirements?

- A. Enable EC2 detailed monitoring, and include network logs. Send all logs through Amazon Kinesis Data Firehose to an Amazon Elasticsearch Service (Amazon ES) cluster that the security team uses for analysis.
- B. Enable VPC Flow Logs for all EC2 instance network interfaces. Publish VPC Flow Logs to an Amazon S3 bucket. Have the security team use Amazon Athena to query and analyze the logs.
- C. Enable access logs for the Application Load Balancer, and publish the logs to an Amazon S3 bucket. Have the security team use Amazon Athena to query and analyze the logs.
- D. Enable Traffic Mirroring and specify all EC2 instance network interfaces as the source.
- E. Send all traffic information through Amazon Kinesis Data Firehose to an Amazon Elasticsearch Service (Amazon ES) cluster that the security team uses for analysis.

Answer: C

Explanation:

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-access-logs.html>

NEW QUESTION 7

- (Exam Topic 1)

An online e-commerce business is running a workload on AWS. The application architecture includes a web tier, an application tier for business logic, and a database tier for user and transactional data management. The database server has a 100 GB memory requirement. The business requires cost-efficient disaster

recovery for the application with an RTO of 5 minutes and an RPO of 1 hour. The business also has a regulatory requirement for out-of-region disaster recovery with a minimum distance between the primary and alternate sites of 250 miles.

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

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

Answer: C

Explanation:

As RTO is in minutes

(<https://docs.aws.amazon.com/wellarchitected/latest/reliability-pillar/plan-for-disaster-recovery-dr.html>) Warm standby (RPO in seconds, RTO in minutes): Maintain a scaled-down version of a fully functional environment always running in the DR Region. Business-critical systems are fully duplicated and are always on, but with a scaled down fleet. When the time comes for recovery, the system is scaled up quickly to handle the production load.

NEW QUESTION 8

- (Exam Topic 1)

A financial services company logs personally identifiable information in its application logs stored in Amazon S3. Due to regulatory compliance requirements, the log files must be encrypted at rest. The security team has mandated that the company's on-premises hardware security modules (HSMs) be used to generate the CMK material.

Which steps should the solutions architect take to meet these requirements?

- A. Create an AWS CloudHSM cluster.
- B. Create a new CMK in AWS KMS using AWS_CloudHSM as the source (or the key material and an origin of AWS_CLOUDHSM).
- C. Enable automatic key rotation on the CMK with a duration of 1 year.
- D. Configure a bucket policy on the logging bucket that disallows uploads of unencrypted data and requires that the encryption source be AWS KMS.
- E. Provision an AWS Direct Connect connection, ensuring there is no overlap of the RFC 1918 address space between on-premises hardware and the VPC.
- F. Configure an AWS bucket policy on the logging bucket that requires all objects to be encrypted.
- G. Configure the logging application to query the on-premises HSMs from the AWS environment for the encryption key material, and create a unique CMK for each logging event.
- H. Create a CMK in AWS KMS with no key material and an origin of EXTERNAL.
- I. Import the key material generated from the on-premises HSMs into the CMK using the public key and import token provided by AWS.
- J. Configure a bucket policy on the logging bucket that disallows uploads of non-encrypted data and requires that the encryption source be AWS KMS.
- K. Create a new CMK in AWS KMS with AWS-provided key material and an origin of AWS_KMS.
- L. Disable this CMK.
- M. and overwrite the key material with the key material from the on-premises HSM using the public key and import token provided by AWS.
- N. Re-enable the CMK.
- O. Enable automatic key rotation on the CMK with a duration of 1 year.
- P. Configure a bucket policy on the logging bucket that disallows uploads of non-encrypted data and requires that the encryption source be AWS KMS.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/security/how-to-byok-bring-your-own-key-to-aws-kms-for-less-than-15-00-a-yea>

<https://docs.aws.amazon.com/kms/latest/developerguide/importing-keys-create-cmk.html>

NEW QUESTION 9

- (Exam Topic 1)

A development team has created a new flight tracker application that provides near-real-time data to users. The application has a front end that consists of an Application Load Balancer (ALB) in front of two large Amazon EC2 instances in a single Availability Zone. Data is stored in a single Amazon RDS MySQL DB instance. An Amazon Route 53 DNS record points to the ALB.

Management wants the development team to improve the solution to achieve maximum reliability with the least amount of operational overhead.

Which set of actions should the team take?

- A. Create RDS MySQL read replica.
- B. Deploy the application to multiple AWS Regions.
- C. Use a Route 53 latency-based routing policy to route to the application.
- D. Configure the DB instance as Multi-AZ.
- E. Deploy the application to two additional EC2 instances in different Availability Zones behind an ALB.
- F. Replace the DB instance with Amazon DynamoDB global table.
- G. Deploy the application in multiple AWS Regions.
- H. Use a Route 53 latency-based routing policy to route to the application.
- I. Replace the DB instance with Amazon Aurora with Aurora Replica.
- J. Deploy the application to multiple smaller EC2 instances across multiple Availability Zones in an Auto Scaling group behind an ALB.

Answer: D

Explanation:

Multi AZ ASG + ALB + Aurora = Less over head and automatic scaling

NEW QUESTION 10

- (Exam Topic 1)

A solutions architect is designing the data storage and retrieval architecture for a new application that a company will be launching soon. The application is designed to ingest millions of small records per minute from devices all around the world. Each record is less than 4 KB in size and needs to be stored in a durable location where it can be retrieved with low latency. The data is ephemeral and the company is required to store the data for 120 days only, after which the data can be deleted.

The solutions architect calculates that, during the course of a year, the storage requirements would be about 10-15 TB.

Which storage strategy is the MOST cost-effective and meets the design requirements?

- A. Design the application to store each incoming record as a single .csv file in an Amazon S3 bucket to allow for indexed retrieval
- B. Configure a lifecycle policy to delete data older than 120 days.
- C. Design the application to store each incoming record in an Amazon DynamoDB table properly configured for the scale
- D. Configure the DynamoDB Time to Live (TTL) feature to delete records older than 120 days.
- E. Design the application to store each incoming record in a single table in an Amazon RDS MySQL database
- F. Run a nightly cron job that executes a query to delete any records older than 120 days.
- G. Design the application to batch incoming records before writing them to an Amazon S3 bucket
- H. Update the metadata for the object to contain the list of records in the batch and use the Amazon S3 metadata search feature to retrieve the data
- I. Configure a lifecycle policy to delete the data after 120 days.

Answer: B

Explanation:

DynamoDB with TTL, cheaper for sustained throughput of small items + suited for fast retrievals. S3 cheaper for storage only, much higher costs with writes. RDS not designed for this use case.

NEW QUESTION 10

- (Exam Topic 1)

A fitness tracking company serves users around the world, with its primary markets in North America and Asia. The company needs to design an infrastructure for its read-heavy user authorization application with the following requirements:

- Be resilient to problems with the application in any Region.
- Write to a database in a single Region.
- Read from multiple Regions.
- Support resiliency across application tiers in each Region.
- Support the relational database semantics reflected in the application. Which combination of steps should a solutions architect take? (Select TWO.)

- A. Use an Amazon Route 53 geoproximity routing policy combined with a multivalue answer routing policy.
- B. Deploy the application, and MySQL database servers to Amazon EC2 instances in each Region
- C. Set up the application so that reads and writes are local to the Region
- D. Create snapshots of the web, application, and database servers and store the snapshots in an Amazon S3 bucket in both Regions
- E. Set up cross-Region replication for the database layer.
- F. Use an Amazon Route 53 geolocation routing policy combined with a failover routing policy.
- G. Set up web, application, and Amazon RDS for MySQL instances in each Region
- H. Set up the application so that reads are local and writes are partitioned based on the user
- I. Set up a Multi-AZ failover for the web, application, and database server
- J. Set up cross-Region replication for the database layer.
- K. Set up active-active web and application servers in each Region
- L. Deploy an Amazon Aurora global database with clusters in each Region
- M. Set up the application to use the in-Region Aurora database endpoint
- N. Create snapshots of the web and application servers and store them in an Amazon S3 bucket in both Regions.

Answer: CE

Explanation:

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html>

Geoproximity routing policy is good to control the user traffic to specific regions. However, a multivalue answer routing policy may cause the users to be randomly sent to other healthy regions that may be far away from the user's location. You can use geolocation routing policy to direct the North American users to your servers on the North America region and configure failover routing to the Asia region in case the North America region fails. You can configure the same for the Asian users pointed to the Asia region servers and have the North America region as its backup.

NEW QUESTION 13

- (Exam Topic 1)

A company has an application that sells tickets online and experiences bursts of demand every 7 days. The application has a stateless presentation layer running on Amazon EC2, an Oracle database to store unstructured data catalog information, and a backend API layer. The front-end layer uses an Elastic Load Balancer to distribute the load across nine On-Demand Instances over three Availability Zones (AZs). The Oracle database is running on a single EC2 instance. The company is experiencing performance issues when running more than two concurrent campaigns. A solutions architect must design a solution that meets the following requirements:

- Address scalability issues.
- Increase the level of concurrency.
- Eliminate licensing costs.
- Improve reliability.

Which set of steps should the solutions architect take?

- A. Create an Auto Scaling group for the front end with a combination of On-Demand and Spot Instances to reduce cost
- B. Convert the Oracle database into a single Amazon RDS reserved DB instance.
- C. Create an Auto Scaling group for the front end with a combination of On-Demand and Spot Instances to reduce cost
- D. Create two additional copies of the database instance, then distribute the databases in separate AZs.
- E. Create an Auto Scaling group for the front end with a combination of On-Demand and Spot Instances to reduce cost

- F. Convert the tables in the Oracle database into Amazon DynamoDB tables.
- G. Convert the On-Demand Instances into Spot Instances to reduce costs for the front end
- H. Convert the tables in the Oracle database into Amazon DynamoDB tables.

Answer: C

Explanation:

Combination of On-Demand and Spot Instances + DynamoDB.

NEW QUESTION 18

- (Exam Topic 1)

A company has developed an application that is running Windows Server on VMware vSphere VMs that the company hosts on-premises. The application data is stored in a proprietary format that must be read through the application. The company manually provisioned the servers and the application.

As part of its disaster recovery plan, the company wants the ability to host its application on AWS temporarily if the company's on-premises environment becomes unavailable. The company wants the application to return to on-premises hosting after a disaster recovery event is complete. The RPO is 15 minutes.

Which solution meets these requirements with the LEAST amount of operational overhead?

- A. Configure AWS DataSync
- B. Replicate the data to Amazon Elastic Block Store (Amazon EBS) volumes. When the on-premises environment is unavailable, use AWS CloudFormation templates to provision Amazon EC2 instances and attach the EBS volumes.
- C. Configure CloudEndure Disaster Recovery. Replicate the data to replication Amazon EC2 instances that are attached to Amazon Elastic Block Store (Amazon EBS) volumes. When the on-premises environment is unavailable, use CloudEndure to launch EC2 instances that use the replicated volumes.
- D. Provision an AWS Storage Gateway. We gateway
- E. Recreate the data to an Amazon S3 bucket
- F. When the on-premises environment is unavailable, use AWS Backup to restore the data to Amazon Elastic Block Store (Amazon EBS) volumes and launch Amazon EC2 instances from these EBS volumes.
- G. Provision an Amazon FSx for Windows File Server file system on AWS. Replicate the data to the file system. When the on-premises environment is unavailable, use AWS CloudFormation templates to provision Amazon EC2 instances and use AWS CloudFormation Init commands to mount the Amazon FSx file shares.

Answer: D

NEW QUESTION 23

- (Exam Topic 1)

A company wants to retire its Oracle Solaris NFS storage arrays. The company requires rapid data migration over its internet network connection to a combination of destinations for Amazon S3, Amazon Elastic File System (Amazon EFS), and Amazon FSx for Windows File Server. The company also requires a full initial copy, as well as incremental transfers of changes until the retirement of the storage arrays. All data must be encrypted and checked for integrity.

What should a solutions architect recommend to meet these requirements?

- A. Configure CloudEndure
- B. Create a project and deploy the CloudEndure agent and token to the storage array
- C. Run the migration plan to start the transfer.
- D. Configure AWS DataSync
- E. Configure the DataSync agent and deploy it to the local network
- F. Create a transfer task and start the transfer.
- G. Configure the aws S3 sync command
- H. Configure the AWS client on the client side with credential
- I. Run the sync command to start the transfer.
- J. Configure AWS Transfer (or FTP)
- K. Configure the FTP client with credential
- L. Script the client to connect and sync to start the transfer.

Answer: B

NEW QUESTION 28

- (Exam Topic 1)

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 data sharing without causing the institute that owns the bucket to assume all the costs for S3 requests and data transfers?

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

Answer: B

Explanation:

In general, bucket owners pay for all Amazon S3 storage and data transfer costs associated with their bucket. A bucket owner, however, can configure a bucket to be a Requester Pays bucket. With Requester Pays buckets, the requester instead of the bucket owner pays the cost of the request and the data download from the bucket. The bucket owner always pays the cost of storing data. If you enable Requester Pays on a bucket, anonymous access to that bucket is not allowed. <https://docs.aws.amazon.com/AmazonS3/latest/userguide/RequesterPaysExamples.html>

NEW QUESTION 30

- (Exam Topic 1)

A company wants to change its internal cloud billing strategy for each of its business units. Currently, the cloud governance team shares reports for overall cloud spending with the head of each business unit. The company uses AWS Organizations to manage the separate AWS accounts for each business unit. The existing tagging standard in Organizations includes the application, environment, and owner. The cloud governance team wants a centralized solution so each business unit receives monthly reports on its cloud spending. The solution should also send notifications for any cloud spending that exceeds a set threshold.

Which solution is the MOST cost-effective way to meet these requirements?

- A. Configure AWS Budgets in each account and configure budget alerts that are grouped by application, environment, and owner
- B. Add each business unit to an Amazon SNS topic for each alert
- C. Use Cost Explorer in each account to create monthly reports for each business unit.
- D. Configure AWS Budgets in the organization's master account and configure budget alerts that are grouped by application, environment, and owner
- E. Add each business unit to an Amazon SNS topic for each alert
- F. Use Cost Explorer in the organization's master account to create monthly reports for each business unit.
- G. Configure AWS Budgets in each account and configure budget alerts that are grouped by application, environment, and owner
- H. Add each business unit to an Amazon SNS topic for each alert
- I. Use the AWS Billing and Cost Management dashboard in each account to create monthly reports for each business unit.
- J. Enable AWS Cost and Usage Reports in the organization's master account and configure reports grouped by application, environment, and owner
- K. Create an AWS Lambda function that processes AWS Cost and Usage Reports, sends budget alerts, and sends monthly reports to each business unit's email list.

Answer: B

Explanation:

Configure AWS Budgets in the organization's master account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in the organization's master account to create monthly reports for each business unit. <https://aws.amazon.com/about-aws/whats-new/2019/07/introducing-aws-budgets-reports/#:~:text=AWS%20Bud>

NEW QUESTION 34

- (Exam Topic 1)

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

Explanation:

Ephemeral ports are not covered in the syllabus so be careful that you don't confuse day to day best practice with what is required for the exam. Link to an explanation on Ephemeral ports here. <https://acloud.guru/forums/aws-certified-solutions-architect-associate/discussion/-KUbcwo4IXefMI7janaK/netw>

NEW QUESTION 36

- (Exam Topic 1)

A solution architect is designing an AWS account structure for a company that consists of multiple terms. All the team will work in the same AWS Region. The company needs a VPC that is connected to the on-premises network. The company expects less than 50 Mbps of total to and from the on-premises network.

Which combination of steps will meet these requirements MOST cost-effectively? (Select TWO)

- A. Create an AWS CloudFormation template that provisions a VPC and the required subnet
- B. Deploy the template to each AWS account
- C. Create an AWS CloudFormation template that provisions a VPC and the required subnet
- D. Deploy the template to a shared services account
- E. Share the subnets by using AWS Resource Access Manager
- F. Use AWS Transit Gateway along with an AWS Site-to-Site VPN for connectivity to the on-premises network
- G. Share the transit gateway by using AWS Resource Access Manager
- H. Use AWS Site-to-Site VPN for connectivity to the on-premises network
- I. Use AWS Direct Connect for connectivity to the on-premises network.

Answer: BD

NEW QUESTION 41

- (Exam Topic 1)

A company has a three-tier application running on AWS with a web server, an application server, and an Amazon RDS MySQL DB instance. A solutions architect is designing a disaster recovery (DR) solution with an RPO of 5 minutes.

Which solution will meet the company's requirements?

- A. Configure AWS Backup to perform cross-Region backups of all servers every 5 minutes

- B. Reprovision the three tiers in the DR Region from the backups using AWS CloudFormation in the event of a disaster.
- C. Maintain another running copy of the web and application server stack in the DR Region using AWS CloudFormation drill detectio
- D. Configure cross-Region snapshots of the DB instance to the DR Region every 5 minute
- E. In the event of a disaster, restore the DB instance using the snapshot in the DR Region.
- F. Use Amazon EC2 Image Builder to create and copy AMIs of the web and application server to both the primary and DR Region
- G. Create a cross-Region read replica of the DB instance in the DR Regio
- H. In the event of a disaster, promote the read replica to become the master and reprovision the servers with AWS CloudFormation using the AMIs.
- I. Create AMIs of the web and application servers in the DR Regio
- J. Use scheduled AWS Glue jobs to synchronize the DB instance with another DB instance in the DR Regio
- K. In the event of a disaster, switch to the DB instance in the DR Region and reprovision the servers with AWS CloudFormation using the AMIs.

Answer: C

Explanation:

deploying a brand new RDS instance will take >30 minutes. You will use EC2 Image builder to put the AMIs into the new region, but not use image builder to LAUNCH them.

NEW QUESTION 43

- (Exam Topic 1)

A finance company is running its business-critical application on current-generation Linux EC2 instances. The application includes a self-managed MySQL database performing heavy I/O operations. The application is working fine to handle a moderate amount of traffic during the month. However, it slows down during the final three days of each month due to month-end reporting, even though the company is using Elastic Load Balancers and Auto Scaling within its infrastructure to meet the increased demand.

Which of the following actions would allow the database to handle the month-end load with the LEAST impact on performance?

- A. Pre-warming Elastic Load Balancers, using a bigger instance type, changing all Amazon EBS volumes to GP2 volumes.
- B. Performing a one-time migration of the database cluster to Amazon RD
- C. and creating several additional read replicas to handle the load during end of month
- D. Using Amazon CioudWatch with AWS Lambda to change the typ
- E. size, or IOPS of Amazon EBS volumes in the cluster based on a specific CloudWatch metric
- F. Replacing all existing Amazon EBS volumes with new PIOPS volumes that have the maximum available storage size and I/O per second by taking snapshots before the end of the month and reverting back afterwards.

Answer: B

Explanation:

In this scenario, the Amazon EC2 instances are in an Auto Scaling group already which means that the database read operations is the possible bottleneck especially during the month-end wherein the reports are generated. This can be solved by creating RDS read replicas.

NEW QUESTION 45

- (Exam Topic 1)

A company is deploying a new cluster for big data analytics on AWS. The cluster will run across many Linux Amazon EC2 instances that are spread across multiple Availability Zones.

All of the nodes in the cluster must have read and write access to common underlying file storage. The file storage must be highly available, must be resilient, must be compatible with the Portable Operating System Interface (POSIX), and must accommodate high levels of throughput.

Which storage solution will meet these requirements?

- A. Provision an AWS Storage Gateway file gateway NFS file share that is attached to an Amazon S3 bucke
- B. Mount the NFS file share on each EC2 instance In the cluster.
- C. Provision a new Amazon Elastic File System (Amazon EFS) file system that uses General Purpose performance mod
- D. Mount the EFS file system on each EC2 instance in the cluster.
- E. Provision a new Amazon Elastic Block Store (Amazon EBS) volume that uses the lo2 volume type. Attach the EBS volume to all of the EC2 instances in the cluster.
- F. Provision a new Amazon Elastic File System (Amazon EFS) file system that uses Max I/O performance mod
- G. Mount the EFS file system on each EC2 instance in the cluster.

Answer: D

NEW QUESTION 47

- (Exam Topic 1)

A company wants to host a new global website that consists of static content. A solutions architect is working on a solution that uses Amazon CloudFront with an origin access identity (OAI) to access website content that is stored in a private Amazon S3 bucket.

During testing, the solutions architect receives 404 errors from the S3 bucket. Error messages appear only for attempts to access paths that end with a forward slash, such as example.com/path/. These requests should return the existing S3 object path/index.html. Any potential solution must not prevent CloudFront from caching the content.

What should the solutions architect do to resolve this problem?

- A. Change the CloudFront origin to an Amazon API Gateway proxy endpoint
- B. Rewrite the S3 request URL by using an AWS Lambda function.
- C. Change the CloudFront origin to an Amazon API Gateway endpoint
- D. Rewrite the S3 request URL in an AWS service integration.
- E. Change the CloudFront configuration to use an AWS Lambda@Edge function that is invoked by a viewer request event to rewrite the S3 request URL.
- F. Change the CloudFront configuration to use an AWS Lambda@Edge function that is invoked by an origin request event to rewrite the S3 request URL.

Answer: C

NEW QUESTION 50

- (Exam Topic 1)

A company is storing data on premises on a Windows file server. The company produces 5 GB of new data daily.

The company migrated part of its Windows-based workload to AWS and needs the data to be available on a file system in the cloud. The company already has established an AWS Direct Connect connection between the on-premises network and AWS. Which data migration strategy should the company use?

- A. Use the file gateway option in AWS Storage Gateway to replace the existing Windows file server, and point the existing file share to the new file gateway.
- B. Use AWS DataSync to schedule a daily task to replicate data between the on-premises Windows file server and Amazon FSx.
- C. Use AWS Data Pipeline to schedule a daily task to replicate data between the on-premises Windows file server and Amazon Elastic File System (Amazon EFS).
- D. Use AWS DataSync to schedule a daily task to replicate data between the on-premises Windows file server and Amazon Elastic File System (Amazon EFS),

Answer: B

Explanation:

<https://aws.amazon.com/storagegateway/file/> <https://docs.aws.amazon.com/fsx/latest/WindowsGuide/migrate-files-to-fsx-datasync.html>
<https://docs.aws.amazon.com/systems-manager/latest/userguide/prereqs-operating-systems.html#prereqs-os-win>

NEW QUESTION 55

- (Exam Topic 1)

A scientific organization requires the processing of text and picture data stored in an Amazon S3 bucket. The data is gathered from numerous radar stations during a mission's live, time-critical phase. The data is uploaded by the radar stations to the source S3 bucket. The data is preceded with the identification number of the radar station.

In a second account, the business built a destination S3 bucket. To satisfy a compliance target, data must be transferred from the source S3 bucket to the destination S3 bucket. Replication is accomplished by using an S3 replication rule that covers all items in the source S3 bucket.

A single radar station has been recognized as having the most precise data. At this radar station, data replication must be completed within 30 minutes of the radar station uploading the items to the source S3 bucket.

What actions should a solutions architect take to ensure that these criteria are met?

- A. Set up an AWS DataSync agent to replicate the prefixed data from the source S3 bucket to the destination S3 bucket
- B. Select to use at available bandwidth on the task, and monitor the task to ensure that it is in the TRANSFERRING status
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert if this status changes.
- D. In the second account, create another S3 bucket to receive data from the radar station with the most accurate data. Set up a new replication rule for this new S3 bucket to separate the replication from the other radar stations. Monitor the maximum replication time to the destination.
- E. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert when the time exceeds the desired threshold.
- F. Enable Amazon S3 Transfer Acceleration on the source S3 bucket, and configure the radar station with the most accurate data to use the new endpoint. Monitor the S3 destination bucket's TotalRequestLatency metric. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert if this status changes.
- G. Create a new S3 replication rule on the source S3 bucket that filters for the keys that use the prefix of the radar station with the most accurate data. Enable S3 Replication Time Control (S3 RTC). Monitor the maximum replication time to the destination. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert when the time exceeds the desired threshold.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-time-control.html>

NEW QUESTION 56

- (Exam Topic 1)

The company needs to determine which costs on the monthly AWS bill are attributable to each application or team. The company also must be able to create reports to compare costs from the last 12 months and to help forecast costs for the next 12 months. A solutions architect must recommend an AWS Billing and Cost Management solution that provides these cost reports.

Which combination of actions will meet these requirements? (Select THREE.)

- A. Activate the user-defined cost allocation tags that represent the application and the team.
- B. Activate the AWS generated cost allocation tags that represent the application and the team.
- C. Create a cost category for each application in Billing and Cost Management.
- D. Activate IAM access to Billing and Cost Management.
- E. Create a cost budget.
- F. Enable Cost Explorer.

Answer: ACF

Explanation:

<https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/manage-cost-categories.html> <https://aws.amazon.com/premiumsupport/knowledge-center/cost-explorer-analyze-spending-and-usage/> <https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/manage-cost-categories.html>
<https://docs.aws.amazon.com/cost-management/latest/userguide/ce-enable.html>

NEW QUESTION 60

- (Exam Topic 1)

A company that tracks medical devices in hospitals wants to migrate its existing storage solution to the AWS Cloud. The company equips all of its devices with sensors that collect location and usage information. This sensor data is sent in unpredictable patterns with large spikes. The data is stored in a MySQL database running on premises at each hospital. The company wants the cloud storage solution to scale with usage.

The company's analytics team uses the sensor data to calculate usage by device type and hospital. The team needs to keep analysis tools running locally while fetching data from the cloud. The team also needs to use existing Java application and SQL queries with as few changes as possible.

How should a solutions architect meet these requirements while ensuring the sensor data is secure?

- A. Store the data in an Amazon Aurora Serverless database.
- B. Serve the data through a Network Load Balancer (NLB). Authenticate users using the NLB with credentials stored in AWS Secrets Manager.
- C. Store the data in an Amazon S3 bucket.
- D. Serve the data through Amazon QuickSight using an IAM user authorized with AWS Identity and Access Management (IAM) with the S3 bucket as the data source.
- E. Store the data in an Amazon Aurora Serverless database.

- F. Serve the data through the Aurora Data API using an IAM user authorized with AWS Identity and Access Management (IAM) and the AWS Secrets Manager ARN.
- G. Store the data in an Amazon S3 bucket
- H. Serve the data through Amazon Athena using AWS PrivateLink to secure the data in transit.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/aws/new-data-api-for-amazon-aurora-serverless/> <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/data-api.html>
<https://aws.amazon.com/blogs/aws/aws-privatelink-for-amazon-s3-now-available/> <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/data-api.html#data-api.access>

The data is currently stored in a MySQL database running on-prem. Storing MySQL data in S3 doesn't sound good so B & D are out. Aurora Data API "enables the SQL HTTP endpoint, a connectionless Web Service API for running SQL queries against this database. When the SQL HTTP endpoint is enabled, you can also query your database from inside the RDS console (these features are free to use)."

NEW QUESTION 65

- (Exam Topic 1)

A company is running a data-intensive application on AWS. The application runs on a cluster of hundreds of Amazon EC2 instances. A shared file system also runs on several EC2 instances that store 200 TB of data. The application reads and modifies the data on the shared file system and generates a report. The job runs once monthly, reads a subset of the files from the shared file system, and takes about 72 hours to complete. The compute instances scale in an Auto Scaling group, but the instances that host the shared file system run continuously. The compute and storage instances are all in the same AWS Region.

A solutions architect needs to reduce costs by replacing the shared file system instances. The file system must provide high performance access to the needed data for the duration of the 72-hour run.

Which solution will provide the LARGEST overall cost reduction while meeting these requirements?

- A. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Intelligent-Tiering storage class
- B. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using lazy loadin
- C. Use the new file system as the shared storage for the duration of the job
- D. Delete the file system when the job is complete.
- E. Migrate the data from the existing shared file system to a large Amazon Elastic Block Store (Amazon EBS) volume with Multi-Attach enable
- F. Attach the EBS volume to each of the instances by using a user data script in the Auto Scaling group launch template
- G. Use the EBS volume as the shared storage for the duration of the job
- H. Detach the EBS volume when the job is complete.
- I. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Standard storage class
- J. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using batch loadin
- K. Use the new file system as the shared storage for the duration of the job
- L. Delete the file system when the job is complete.
- M. Migrate the data from the existing shared file system to an Amazon S3 bucket
- N. Before the job runs each month, use AWS Storage Gateway to create a file gateway with the data from Amazon S3. Use the file gateway as the shared storage for the job
- O. Delete the file gateway when the job is complete.

Answer: B

NEW QUESTION 68

- (Exam Topic 1)

A company has application services that have been containerized and deployed on multiple Amazon EC2 instances with public IPs. An Apache Kafka cluster has been deployed to the EC2 instances. A PostgreSQL database has been migrated to Amazon RDS for PostgreSQL. The company expects a significant increase of orders on its platform when a new version of its flagship product is released.

What changes to the current architecture will reduce operational overhead and support the product release?

- A. Create an EC2 Auto Scaling group behind an Application Load Balance
- B. Create additional read replicas for the DB instance
- C. Create Amazon Kinesis data streams and configure the application services to use the data stream
- D. Store and serve static content directly from Amazon S3.
- E. Create an EC2 Auto Scaling group behind an Application Load Balance
- F. Deploy the DB instance in Multi-AZ mode and enable storage auto scaling
- G. Create Amazon Kinesis data streams and configure the application services to use the data stream
- H. Store and serve static content directly from Amazon S3.
- I. Deploy the application on a Kubernetes cluster created on the EC2 instances behind an Application Load Balance
- J. Deploy the DB instance in Multi-AZ mode and enable storage auto scaling
- K. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluster
- L. Store static content in Amazon S3 behind an Amazon CloudFront distribution.
- M. Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate and enable auto scaling behind an Application Load Balance
- N. Create additional read replicas for the DB instance
- O. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluster
- P. Store static content in Amazon S3 behind an Amazon CloudFront distribution.

Answer: D

Explanation:

Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate and enable auto scaling behind an Application Load Balancer. Create additional read replicas for the DB instance. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluster. Store static content in Amazon S3 behind an Amazon CloudFront distribution.

NEW QUESTION 73

- (Exam Topic 1)

A solutions architect needs to advise a company on how to migrate its on-premises data processing application to the AWS Cloud. Currently, users upload input files through a web portal. The web server then stores the uploaded files on NAS and messages the processing server over a message queue. Each media file can take up to 1 hour to process. The company has determined that the number of media files awaiting processing is significantly higher during business hours, with

the number of files rapidly declining after business hours.
What is the MOST cost-effective migration recommendation?

- A. Create a queue using Amazon SQ
- B. Configure the existing web server to publish to the new queue. When there are messages in the queue, invoke an AWS Lambda function to pull requests from the queue and process the file
- C. Store the processed files in an Amazon S3 bucket.
- D. Create a queue using Amazon M
- E. Configure the existing web server to publish to the new queue. When there are messages in the queue, create a new Amazon EC2 instance to pull requests from the queue and process the file
- F. Store the processed files in Amazon EF
- G. Shut down the EC2 instance after the task is complete.
- H. Create a queue using Amazon M
- I. Configure the existing web server to publish to the new queue. When there are messages in the queue, invoke an AWS Lambda function to pull requests from the queue and process the file
- J. Store the processed files in Amazon EFS.
- K. Create a queue using Amazon SO
- L. Configure the existing web server to publish to the new queu
- M. Use Amazon EC2 instances in an EC2 Auto Scaling group to pull requests from the queue and process the file
- N. Scale the EC2 instances based on the SOS queue lengt
- O. Store the processed files in an Amazon S3 bucket.

Answer: D

Explanation:

<https://aws.amazon.com/blogs/compute/operating-lambda-performance-optimization-part-1/>

NEW QUESTION 74

- (Exam Topic 1)

A financial services company receives a regular data feed from its credit card servicing partner. Approximately 5.1 records are sent every 15 minutes in plaintext, delivered over HTTPS directly into an Amazon S3 bucket with server-side encryption. This feed contains sensitive credit card primary account number (PAN) data. The company needs to automatically mask the PAN before sending the data to another S3 bucket for additional internal processing. The company also needs to remove and merge specific fields, and then transform the record into JSON format. Additionally, extra feeds are likely to be added in the future, so any design needs to be easily expandable.

Which solutions will meet these requirements?

- A. Trigger an AWS Lambda function on file delivery that extracts each record and writes it to an Amazon SQS queue
- B. Trigger another Lambda function when new messages arrive in the SOS queue to process the records, writing the results to a temporary location in Amazon S3. Trigger a final Lambda function once the SOS queue is empty to transform the records into JSON format and send the results to another S3 bucket for internal processing.
- C. Trigger an AWS Lambda function on file delivery that extracts each record and writes it to an Amazon SOS queue
- D. Configure an AWS Fargate container application to
- E. automatically scale to a single instance when the SOS queue contains message
- F. Have the application process each record, and transform the record into JSON format
- G. When the queue is empty, send the results to another S3 bucket for internal processing and scale down the AWS Fargate instance.
- H. Create an AWS Glue crawler and custom classifier based on the data feed formats and build a table definition to match. Trigger an AWS Lambda function on file delivery to start an AWS Glue ETL job to transform the entire record according to the processing and transformation requirements
- I. Define the output format as JSON
- J. Once complete, have the ETL job send the results to another S3 bucket for internal processing.
- K. Create an AWS Glue crawler and custom classifier based upon the data feed formats and build a table definition to match
- L. Perform an Amazon Athena query on file delivery to start an Amazon EMR ETL job to transform the entire record according to the processing and transformation requirements
- M. Define the output format as JSON
- N. Once complete, send the results to another S3 bucket for internal processing and scale down the EMR cluster.

Answer: C

Explanation:

You can use a Glue crawler to populate the AWS Glue Data Catalog with tables. The Lambda function can be triggered using S3 event notifications when object create events occur. The Lambda function will then trigger the Glue ETL job to transform the records, masking the sensitive data and modifying the output format to JSON. This solution meets all requirements.

Create an AWS Glue crawler and custom classifier based on the data feed formats and build a table definition to match. Trigger an AWS Lambda function on file delivery to start an AWS Glue ETL job to transform the entire record according to the processing and transformation requirements. Define the output format as JSON. Once complete, have the ETL job send the results to another S3 bucket for internal processing.

<https://docs.aws.amazon.com/glue/latest/dg/trigger-job.html>

https://d1.awsstatic.com/Products/product-name/diagrams/product-page-diagram_Glue_Event-driven-ETL-Pipel

NEW QUESTION 77

- (Exam Topic 1)

A company is running an Apache Hadoop cluster on Amazon EC2 instances. The Hadoop cluster stores approximately 100 TB of data for weekly operational reports and allows occasional access for data scientists to retrieve data. The company needs to reduce the cost and operational complexity for storing and serving this data.

Which solution meets these requirements in the MOST cost-effective manner?

- A. Move the Hadoop cluster from EC2 instances to Amazon EM
- B. Allow data access patterns to remain the same.
- C. Write a script that resizes the EC2 instances to a smaller instance type during downtime and resizes the instances to a larger instance type before the reports are created.
- D. Move the data to Amazon S3 and use Amazon Athena to query the data for reports
- E. Allow the data scientists to access the data directly in Amazon S3.
- F. Migrate the data to Amazon DynamoDB and modify the reports to fetch data from DynamoDB

G. Allow the data scientists to access the data directly in DynamoDB.

Answer: C

Explanation:

"The company needs to reduce the cost and operational complexity for storing and serving this data. Which solution meets these requirements in the MOST cost-effective manner?" EMR storage is ephemeral. The company has 100TB that need to persist, they would have to use EMRFS to backup to S3 anyway.

<https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-plan-storage.html>

100TB

EBS - 8.109\$ S3 - 2.355\$

You have saved 5.752\$

This amount can be used for Athen. BTW. we don't know indexes, amount of data that is scanned. What we know is that it will be: "occasional access for data scientists to retrieve data"

NEW QUESTION 78

- (Exam Topic 1)

A company needs to store and process image data that will be uploaded from mobile devices using a custom mobile app. Usage peaks between 8 AM and 5 PM on weekdays, with thousands of uploads per minute. The app is rarely used at any other time A user is notified when image processing is complete.

Which combination of actions should a solutions architect take to ensure image processing can scale to handle the load? (Select THREE.)

- A. Upload files from the mobile software directly to Amazon S3. Use S3 event notifications to create a message in an Amazon MQ queue.
- B. Upload files from the mobile software directly to Amazon S3. Use S3 event notifications to create a message in an Amazon Simple Queue Service (Amazon SQS) standard queue.
- C. Invoke an AWS Lambda function to perform image processing when a message is available in the queue.
- D. Invoke an S3 Batch Operations job to perform image processing when a message is available in the queue.
- E. Send a push notification to the mobile app by using Amazon Simple Notification Service (Amazon SNS) when processing is complete.
- F. Send a push notification to the mobile app by using Amazon Simple Email Service (Amazon SES) when processing is complete.

Answer: BCE

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/batch-ops-basics.html>

NEW QUESTION 83

- (Exam Topic 1)

A company has a photo sharing social networking application. To provide a consistent experience for users, the company performs some image processing on the photos uploaded by users before publishing on the application. The image processing is implemented using a set of Python libraries.

The current architecture is as follows:

- The image processing Python code runs in a single Amazon EC2 instance and stores the processed images in an Amazon S3 bucket named ImageBucket.
- The front-end application, hosted in another bucket, loads the images from ImageBucket to display to users. With plans for global expansion, the company wants to implement changes in its existing architecture to be able to scale for increased demand on the application and reduce management complexity as the application scales.

Which combination of changes should a solutions architect make? (Select TWO.)

- A. Place the image processing EC2 instance into an Auto Scaling group.
- B. Use AWS Lambda to run the image processing tasks.
- C. Use Amazon Rekognition for image processing.
- D. Use Amazon CloudFront in front of ImageBucket.
- E. Deploy the applications in an Amazon ECS cluster and apply Service Auto Scaling.

Answer: BD

Explanation:

<https://prismatic.io/blog/why-we-moved-from-lambda-to-ecs/>

NEW QUESTION 87

- (Exam Topic 1)

A media company uses Amazon DynamoDB to store metadata for its catalog of movies that are available to stream. Each media item Contains user-facing content that concludes a description of the media, a list of search tags, and similar data. In addition, media items include a list of Amazon S3 key names that relate to movie files. The company stores these movie files in a single S3 bucket that has versioning enable. The company uses Amazon CloudFront to serve these movie files.

The company has 100.000 media items, and each media item can have many different S3 objects that represent different encodings of the same media S3 objects that belong to the same media item are grouped together under the same key prefix, which is a random unique ID

Because of an expiring contract with a media provider, the company must remove 2.000 media Items. The company must completely delete all DynamoDB keys and movie files on Amazon S3 that are related to these media items within 36 hours The company must ensure that the content cannot be recovered.

Which combination of actions will meet these requirements? (Select TWO.)

- A. Configure the dynamoDB table with a TTL field
- B. Create and invoke an AWS Lambda function to perform a conditional update Set the TTL field to the time of the contract's expiration on every affected media item.
- C. Configure an S3 Lifecycle object expiration rule that is based on the contract's expiration date
- D. Write a script to perform a conditional delete on all the affected DynamoDB records
- E. Temporarily suspend versioning on the S3 bucket
- F. Create and invoke an AWS Lambda function that deletes affected objects Reactivate versioning when the operation is complete
- G. Write a script to delete objects from Amazon S3 Specify in each request a NoncurrentVersionExpiration property with a NoncurrentDays attribute set to 0.

Answer: CE

NEW QUESTION 90

- (Exam Topic 1)

A company built an ecommerce website on AWS using a three-tier web architecture. The application is Java-based and composed of an Amazon CloudFront distribution, an Apache web server layer of Amazon EC2 instances in an Auto Scaling group, and a backend Amazon Aurora MySQL database.

Last month, during a promotional sales event, users reported errors and timeouts while adding items to their shopping carts. The operations team recovered the logs created by the web servers and reviewed Aurora DB cluster performance metrics. Some of the web servers were terminated before logs could be collected and the Aurora metrics were not sufficient for query performance analysis.

Which combination of steps must the solutions architect take to improve application performance visibility during peak traffic events? (Select THREE.)

- A. Configure the Aurora MySQL DB cluster to publish slow query and error logs to Amazon CloudWatch Logs.
- B. Implement the AWS X-Ray SDK to trace incoming HTTP requests on the EC2 instances and implement tracing of SQL queries with the X-Ray SDK for Java.
- C. Configure the Aurora MySQL DB cluster to stream slow query and error logs to Amazon Kinesis.
- D. Install and configure an Amazon CloudWatch Logs agent on the EC2 instances to send the Apache logs to CloudWatch Logs.
- E. Enable and configure AWS CloudTrail to collect and analyze application activity from Amazon EC2 and Aurora.
- F. Enable Aurora MySQL DB cluster performance benchmarking and publish the stream to AWS X-Ray.

Answer: ABD

Explanation:

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_LogAccess.Concepts.MySQL.html <https://aws.amazon.com/blogs/mt/simplifying-apache-server-logs-with-amazon-cloudwatch-logs-insights/> <https://docs.aws.amazon.com/xray/latest/devguide/xray-sdk-dotnet-messagehandler.html>
<https://docs.aws.amazon.com/xray/latest/devguide/xray-sdk-java-sqlclients.html>

NEW QUESTION 93

- (Exam Topic 1)

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 connection.
- 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 connection.
- 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 connection.
- G. Use Direct Connect Gateway to access data in other AWS Regions.

Answer: D

Explanation:

This feature also allows you to connect to any of the participating VPCs from any Direct Connect location, further reducing your costs for making using AWS services on a cross-region basis. <https://aws.amazon.com/blogs/aws/new-aws-direct-connect-gateway-inter-region-vpc-access/>
<https://docs.aws.amazon.com/whitepapers/latest/aws-vpc-connectivity-options/aws-direct-connect-aws-transit-g>

NEW QUESTION 94

- (Exam Topic 1)

A company is building a hybrid solution between its existing on-premises systems and a new backend in AWS. The company has a management application to monitor the state of its current IT infrastructure and automate responses to issues. The company wants to incorporate the status of its consumed AWS services into the application. The application uses an HTTPS endpoint to receive updates.

Which approach meets these requirements with the LEAST amount of operational overhead?

- A. Configure AWS Systems Manager OpsCenter to ingest operational events from the on-premises systems. Retire the on-premises management application and adopt OpsCenter as the hub.
- B. Configure Amazon EventBridge (Amazon CloudWatch Events) to detect and react to changes for AWS Health events from the AWS Personal Health Dashboard. Configure the EventBridge (CloudWatch Events) event to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic and subscribe the topic to the HTTPS endpoint of the management application.
- C. Modify the on-premises management application to call the AWS Health API to poll for status events of AWS services.
- D. Configure Amazon EventBridge (Amazon CloudWatch Events) to detect and react to changes for AWS Health events from the AWS Service Health Dashboard. Configure the EventBridge (CloudWatch Events) event to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic and subscribe the topic to an HTTPS endpoint for the management application with a topic filter corresponding to the services being used.

Answer: A

Explanation:

ALB & NLB both supports IPs as targets. Questions is based on TCP traffic over VPN to on-premise. TCP is layer 4 and the , load balancer should be NLB. Then next questions does NLB supports loadbalancing traffic over VPN. And answer is YES based on below URL.

<https://aws.amazon.com/about-aws/whats-new/2018/09/network-load-balancer-now-supports-aws-vpn/>

Target as IPs for NLB & ALB: <https://aws.amazon.com/elasticloadbalancing/faqs/?nc=sn&loc=5> <https://aws.amazon.com/elasticloadbalancing/application-load-balancer/>

NEW QUESTION 96

- (Exam Topic 1)

A solutions architect is responsible for redesigning a legacy Java application to improve its availability, data durability, and scalability. Currently, the application runs on a single high-memory Amazon EC2 instance. It accepts HTTP requests from upstream clients, adds them to an in-memory queue, and responds with a 200 status. A separate application thread reads items from the queue, processes them, and persists the results to an Amazon RDS MySQL instance. The processing time for each item takes 90 seconds on average, most of which is spent waiting on external service calls, but the application is written to process multiple items in parallel.

Traffic to this service is unpredictable. During periods of high load, items may sit in the internal queue for over an hour while the application processes the backlog.

In addition, the current system has issues with availability and data loss if the single application node fails. Clients that access this service cannot be modified. They expect to receive a response to each HTTP request they send within 10 seconds before they will time out and retry the request.

Which approach would improve the availability and durability of the system while decreasing the processing latency and minimizing costs?

- A. Create an Amazon API Gateway REST API that uses Lambda proxy integration to pass requests to an AWS Lambda function
- B. Migrate the core processing code to a Lambda function and write a wrapper class that provides a handler method that converts the proxy events to the internal application data model and invokes the processing module.
- C. Create an Amazon API Gateway REST API that uses a service proxy to put items in an Amazon SQS queue
- D. Extract the core processing code from the existing application and update it to pull items from Amazon SQS instead of an in-memory queue
- E. Deploy the new processing application to smaller EC2 instances within an Auto Scaling group that scales dynamically based on the approximate number of messages in the Amazon SQS queue.
- F. Modify the application to use Amazon DynamoDB instead of Amazon RDS
- G. Configure Auto Scaling for the DynamoDB table
- H. Deploy the application within an Auto Scaling group with a scaling policy based on CPU utilization
- I. Back the in-memory queue with a memory-mapped file to an instance store volume and periodically write that file to Amazon S3.
- J. Update the application to use a Redis task queue instead of the in-memory queue
- K. Build a Docker container image for the application
- L. Create an Amazon ECS task definition that includes the application container and a separate container to host Redis
- M. Deploy the new task definition as an ECS service using AWS Fargate, and enable Auto Scaling.

Answer: B

Explanation:

The obvious challenges here are long workloads, scalability based on queue load, and reliability. Almost always the default answer to queue related workload is SQS. Since the workloads are very long (90 minutes) Lambdas cannot be used (15 mins max timeout). So, autoscaled smaller EC2 nodes that wait on external services to complete the task makes more sense. If the task fails, the message is returned to the queue and retried.

NEW QUESTION 100

- (Exam Topic 1)

A company plans to migrate to AWS. A solutions architect uses AWS Application Discovery Service over the fleet and discovers that there is an Oracle data warehouse and several PostgreSQL databases. Which combination of migration patterns will reduce licensing costs and operational overhead? (Select TWO.)

- A. Lift and shift the Oracle data warehouse to Amazon EC2 using AWS DMS.
- B. Migrate the Oracle data warehouse to Amazon Redshift using AWS SCT and AWS QMS.
- C. Lift and shift the PostgreSQL databases to Amazon EC2 using AWS DMS.
- D. Migrate the PostgreSQL databases to Amazon RDS for PostgreSQL using AWS DMS
- E. Migrate the Oracle data warehouse to an Amazon EMR managed cluster using AWS DMS.

Answer: BD

Explanation:

<https://aws.amazon.com/getting-started/hands-on/migrate-oracle-to-amazon-redshift/> <https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/migrate-an-on-premises-postgresql-database>

NEW QUESTION 102

- (Exam Topic 1)

A large company is running a popular web application. The application runs on several Amazon EC2 Linux Instances in an Auto Scaling group in a private subnet. An Application Load Balancer is targeting the Instances in the Auto Scaling group in the private subnet. AWS Systems Manager Session Manager is configured, and AWS Systems Manager Agent is running on all the EC2 instances.

The company recently released a new version of the application. Some EC2 instances are now being marked as unhealthy and are being terminated. As a result, the application is running at reduced capacity. A solutions architect tries to determine the root cause by analyzing Amazon CloudWatch logs that are collected from the application, but the logs are inconclusive.

How should the solutions architect gain access to an EC2 instance to troubleshoot the issue?

- A. Suspend the Auto Scaling group's HealthCheck scaling process
- B. Use Session Manager to log in to an instance that is marked as unhealthy
- C. Enable EC2 instance termination protection. Use Session Manager to log in to an instance that is marked as unhealthy.
- D. Set the termination policy to OldestInstance on the Auto Scaling group
- E. Use Session Manager to log in to an instance that is marked as unhealthy
- F. Suspend the Auto Scaling group's Terminate process
- G. Use Session Manager to log in to an instance that is marked as unhealthy

Answer: D

Explanation:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-suspend-resume-processes.html>

It shows that for Amazon EC2 Auto Scaling, there are two primary process types: Launch and Terminate. The Launch process adds a new Amazon EC2 instance to an Auto Scaling group, increasing its capacity. The Terminate process removes an Amazon EC2 instance from the group, decreasing its capacity. HealthCheck process for EC2 autoscaling is not a primary process! It is a process along with the following AddToLoadBalancer AlarmNotification AZRebalance HealthCheck InstanceRefresh ReplaceUnhealthy ScheduledActions. From the requirements, some EC2 instances are now being marked as unhealthy and are being terminated. Application is running at reduced capacity not because instances are marked unhealthy but because they are being terminated.

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-suspend-resume-processes.html#choosing-suspend-r>

NEW QUESTION 104

- (Exam Topic 1)

A solutions architect is evaluating the reliability of a recently migrated application running on AWS. The front end is hosted on Amazon S3 and accelerated by Amazon CloudFront. The application layer is running in a stateless Docker container on an Amazon EC2 On-Demand Instance with an Elastic IP address. The storage layer is a MongoDB database running on an EC2 Reserved Instance in the same Availability Zone as the application layer.

Which combination of steps should the solutions architect take to eliminate single points of failure with minimal application code changes? (Select TWO.)

- A. Create a REST API in Amazon API Gateway and use AWS Lambda functions as the application layer.
- B. Create an Application Load Balancer and migrate the Docker container to AWS Fargate.
- C. Migrate the storage layer to Amazon DynamoD8.
- D. Migrate the storage layer to Amazon DocumentD8 (with MongoDB compatibility).
- E. Create an Application Load Balancer and move the storage layer to an EC2 Auto Scaling group.

Answer: BD

Explanation:

https://aws.amazon.com/documentdb/?nc1=h_ls

<https://aws.amazon.com/blogs/containers/using-alb-ingress-controller-with-amazon-eks-on-fargate/>

NEW QUESTION 107

- (Exam Topic 1)

A company runs an application that gives users the ability to search for videos and related information by using keywords that are curated from content providers. The application data is stored in an on-premises Oracle database that is 800 GB in size.

The company wants to migrate the data to an Amazon Aurora MySQL DB instance. A solutions architect plans to use the AWS Schema Conversion Tool and AWS Database Migration Service (AWS DMS) for the migration. During the migration, the existing database must serve ongoing requests. The migration must be completed with minimum downtime. Which solution will meet these requirements?

- A. Create primary key indexes, secondary indexes, and referential integrity constraints in the target database before starting the migration process
- B. Use AWS DMS to run the conversion report for Oracle to Aurora MySQL
- C. Remediate any issues. Then use AWS DMS to migrate the data
- D. Use the M5 or CS DMS replication instance type for ongoing replication
- E. Turn off automatic backups and logging of the target database until the migration and cutover processes are complete

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Managing.Backups.html>

NEW QUESTION 109

- (Exam Topic 1)

A company provides a centralized Amazon EC2 application hosted in a single shared VPC. The centralized application must be accessible from client applications running in the VPCs of other business units. The centralized application front end is configured with a Network Load Balancer (NLB) for scalability.

Up to 10 business unit VPCs will need to be connected to the shared VPC. Some of the business unit VPC CIDR blocks overlap with the shared VPC, and some overlap with each other. Network connectivity to the centralized application in the shared VPC should be allowed from authorized business unit VPCs only. Which network configuration should a solutions architect use to provide connectivity from the client applications in the business unit VPCs to the centralized application in the shared VPC?

- A. Create an AWS Transit Gateway
- B. Attach the shared VPC and the authorized business unit VPCs to the transit gateway
- C. Create a single transit gateway route table and associate it with all of the attached VPC
- D. Allow automatic propagation of routes from the attachments into the route table
- E. Configure VPC routing tables to send traffic to the transit gateway.
- F. Create a VPC endpoint service using the centralized application NLB and enable (the option to require endpoint acceptance)
- G. Create a VPC endpoint in each of the business unit VPCs using the service name of the endpoint service
- H. Accept authorized endpoint requests from the endpoint service console.
- I. Create a VPC peering connection from each business unit VPC to the shared VPC
- J. Accept the VPC peering connections from the shared VPC console
- K. Configure VPC routing tables to send traffic to the VPC peering connection.
- L. Configure a virtual private gateway for the shared VPC and create customer gateways for each of the authorized business unit VPC
- M. Establish a Site-to-Site VPN connection from the business unit VPCs to the shared VPC
- N. Configure VPC routing tables to send traffic to the VPN connection.

Answer: B

Explanation:

Amazon Transit Gateway doesn't support routing between Amazon VPCs with overlapping CIDRs. If you attach a new Amazon VPC that has a CIDR which overlaps with an already attached Amazon VPC, Amazon Transit Gateway will not propagate the new Amazon VPC route into the Amazon Transit Gateway route table.

<https://docs.aws.amazon.com/elasticloadbalancing/latest/network/load-balancer-target-groups.html#client-ip-pre>

NEW QUESTION 111

- (Exam Topic 1)

A company is launching a new web application on Amazon EC2 instances. Development and production workloads exist in separate AWS accounts.

According to the company's security requirements, only automated configuration tools are allowed to access the production account. The company's security team wants to receive immediate notification if any manual access to the production AWS account or EC2 instances occurs.

Which combination of actions should a solutions architect take in the production account to meet these requirements? (Select THREE.)

- A. Turn on AWS CloudTrail logs in the application's primary AWS Region. Use Amazon Athena to query the logs for AwsConsoleSignIn events.
- B. Configure Amazon Simple Email Service (Amazon SES) to send email to the security team when an alarm is activated.
- C. Deploy EC2 instances in an Auto Scaling group. Configure the launch template to deploy instances without key pairs. Configure Amazon CloudWatch Logs to capture system access logs. Create an Amazon CloudWatch alarm that is based on the logs to detect when a user logs in to an EC2 instance.
- D. Configure an Amazon Simple Notification Service (Amazon SNS) topic to send a message to the security team when an alarm is activated.
- E. Turn on AWS CloudTrail logs for all AWS Region.
- F. Configure Amazon CloudWatch alarms to provide an alert when an AwsConsoleSignIn event is detected.
- G. Deploy EC2 instances in an Auto Scaling group.

- H. Configure the launch template to delete the key pair after launch
I. Configure Amazon CloudWatch Logs for the system access logs Create an Amazon CloudWatch dashboard to show user logins over time.

Answer: CDE

NEW QUESTION 116

- (Exam Topic 1)

A company is running an application distributed over several Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer. The security team requires that all application access attempts be made available for analysis. Information about the client IP address, connection type, and user agent must be included.

Which solution will meet these requirements?

- A. Enable EC2 detailed monitoring, and include network log
- B. Send all logs through Amazon Kinesis Data Firehose to an Amazon Elasticsearch Service (Amazon ES) cluster that the security team uses for analysis.
- C. Enable VPC Flow Logs for all EC2 instance network interfaces. Publish VPC Flow Logs to an Amazon S3 bucket. Have the security team use Amazon Athena to query and analyze the logs.
- D. Enable access logs for the Application Load Balancer, and publish the logs to an Amazon S3 bucket. Have the security team use Amazon Athena to query and analyze the logs.
- E. Enable Traffic Mirroring and specify all EC2 instance network interfaces as the source.
- F. Send all traffic information through Amazon Kinesis Data Firehose to an Amazon Elasticsearch Service (Amazon ES) cluster that the security team uses for analysis.

Answer: C

Explanation:

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-access-logs.html> <https://docs.aws.amazon.com/vpc/latest/mirroring/what-is-traffic-mirroring.html>

NEW QUESTION 118

- (Exam Topic 1)

A solutions architect works for a government agency that has strict disaster recovery requirements. All Amazon Elastic Block Store (Amazon EBS) snapshots are required to be saved in at least two additional AWS Regions. The agency also is required to maintain the lowest possible operational overhead.

Which solution meets these requirements?

- A. Configure a policy in Amazon Data Lifecycle Manager (Amazon DLM) to run once daily to copy the EBS snapshots to the additional Regions.
- B. Use Amazon EventBridge (Amazon CloudWatch Events) to schedule an AWS Lambda function to copy the EBS snapshots to the additional Regions.
- C. Set up AWS Backup to create the EBS snapshot.
- D. Configure Amazon S3 cross-Region replication to copy the EBS snapshots to the additional Regions.
- E. Schedule Amazon EC2 Image Builder to run once daily to create an AMI and copy the AMI to the additional Regions.

Answer: B

NEW QUESTION 120

- (Exam Topic 1)

A company manages an on-premises JavaScript front-end web application. The application is hosted on two servers secured with a corporate Active Directory. The application calls a set of Java-based microservices on an application server and stores data in a clustered MySQL database. The application is heavily used during the day on weekdays. It is lightly used during the evenings and weekends.

Daytime traffic to the application has increased rapidly, and reliability has diminished as a result. The company wants to migrate the application to AWS with a solution that eliminates the need for server maintenance, with an API to securely connect to the microservices.

Which combination of actions will meet these requirements? (Select THREE.)

- A. Host the web application on Amazon S3. Use Amazon Cognito identity pools (federated identities) with SAML for authentication and authorization.
- B. Host the web application on Amazon EC2 with Auto Scaling.
- C. Use Amazon Cognito federation and Login with Amazon for authentication and authorization.
- D. Create an API layer with Amazon API Gateway.
- E. Rehost the microservices on AWS Fargate containers.
- F. Create an API layer with Amazon API Gateway.
- G. Rehost the microservices on Amazon Elastic Container Service (Amazon ECS) containers.
- H. Replatform the database to Amazon RDS for MySQL.
- I. Replatform the database to Amazon Aurora MySQL Serverless.

Answer: ACE

NEW QUESTION 121

- (Exam Topic 1)

A company with global offices has a single 1 Gbps AWS Direct Connect connection to a single AWS Region. The company's on-premises network uses the connection to communicate with the company's resources in the AWS Cloud. The connection has a single private virtual interface that connects to a single VPC.

A solutions architect must implement a solution that adds a redundant Direct Connect connection in the same Region. The solution also must provide connectivity to other Regions through the same pair of Direct Connect connections as the company expands into other Regions.

Which solution meets these requirements?

- A. Provision a Direct Connect gateway.
- B. Delete the existing private virtual interface from the existing connection.
- C. Create the second Direct Connect connection.
- D. Create a new private virtual interface on each connection, and connect both private virtual interfaces to the Direct Connect gateway.
- E. Connect the Direct Connect gateway to the single VPC.
- F. Keep the existing private virtual interface.
- G. Create the second Direct Connect connection.
- H. Create a new private virtual interface on the new connection, and connect the new private virtual interface to the single VPC.

- I. Keep the existing private virtual interface
- J. Create the second Direct Connect connectio
- K. Create a new public virtual interface on the new connection, and connect the new public virtual interface to the single VPC.
- L. Provision a transit gatewa
- M. Delete the existing private virtual interface from the existing connection.Create the second Direct Connect connectio
- N. Create a new private virtual interface on each connection, and connect both private virtual interfaces to the transit gatewa
- O. Associate the transit gateway with the single VPC.

Answer: A

Explanation:

A Direct Connect gateway is a globally available resource. You can create the Direct Connect gateway in any Region and access it from all other Regions. The following describe scenarios where you can use a Direct Connect gateway.

<https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-gateways-intro.html>

NEW QUESTION 122

- (Exam Topic 1)

A medical company is running a REST API on a set of Amazon EC2 instances. The EC2 instances run in an Auto Scaling group behind an Application Load Balancer (ALB). The ALB runs in three public subnets, and the EC2 instances run in three private subnets. The company has deployed an Amazon CloudFront distribution that has the ALB as the only origin.

Which solution should a solutions architect recommend to enhance the origin security?

- A. Store a random string in AWS Secrets Manage
- B. Create an AWS Lambda (unction for automatic secret rotatio
- C. Configure CloudFront to inject the random string as a custom HTTP header for the origin reques
- D. Create an AWS WAF web ACL rule with a string match rule for the custom heade
- E. Associate the web ACL with the ALB.
- F. Create an AWS WAF web ACL rule with an IP match condition of the CloudFront service IP address range
- G. Associate the web ACL with the AL
- H. Move the ALB into the three private subnets.
- I. Store a random string in AWS Systems Manager Parameter Stor
- J. Configure Parameter Store automatic rotation for the strin
- K. Configure CloudFront to inject the random siring as a custom HTTP header for the origin reques
- L. Inspect the value of the custom HTTP header, and block access in the ALB.
- M. Configure AWS Shield Advance
- N. Create a security group policy to allow connections from CloudFront service IP address range
- O. Add the policy to AWS Shield Advanced, and attach the policy to the ALB.

Answer: D

Explanation:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-suspend-resume-processes.html>

it shows For Amazon EC2 Auto Scaling, there are two primary process types: Launch and Terminate. The Launch process adds a new Amazon EC2 instance to an Auto Scaling group, increasing its capacity. The Terminate process removes an Amazon EC2 instance from the group, decreasing its capacity. HealthCheck process for EC2 autoscaling is not a primary process! It is a process along with the following AddToLoadBalancer AlarmNotification AZRebalance HealthCheck InstanceRefresh ReplaceUnhealthy ScheduledActions From the requirements, Some EC2 instances are now being marked as unhealthy and are being terminated. Application is running at reduced capacity not because instances are marked unhealthy but because they are being terminated.

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-suspend-resume-processes.html#choosing-suspend-r>

NEW QUESTION 127

- (Exam Topic 1)

A company hosts a web application that runs on a group of Amazon EC2 instances that are behind an Application Load Balancer (ALB) in a VPC. The company wants to analyze the network payloads to reverse-engineer a sophisticated attack of the application.

Which approach should the company take to achieve this goal?

- A. Enable VPC Flow Log
- B. Store the flow logs in an Amazon S3 bucket for analysis.
- C. Enable Traffic Mirroring on the network interface of the EC2 instance
- D. Send the mirrored traffic to a target for storage and analysis.
- E. Create an AWS WAF web AC
- F. and associate it with the AL
- G. Configure AWS WAF logging.
- H. Enable logging for the AL
- I. Store the logs in an Amazon S3 bucket for analysis.

Answer: A

NEW QUESTION 131

- (Exam Topic 1)

A company wants to migrate an application to Amazon EC2 from VMware Infrastructure that runs in an on-premises data center. A solutions architect must preserve the software and configuration settings during the migration.

What should the solutions architect do to meet these requirements?

- A. Configure the AWS DataSync agent to start replicating the data store to Amazon FSx for Windows File Server Use the SMB share to host the VMware data stor
- B. Use VM Import/Export to move the VMs to Amazon EC2.
- C. Use the VMware vSphere client to export the application as an image in Open Virealization Format (OVF) format Create an Amazon S3 bucket to store the image in the destination AWS Regio
- D. Create and apply an IAM role for VM Import Use the AWS CLI to run the EC2 import command.
- E. Configure AWS Storage Gateway for files service to export a Common Internet File System (CIFS) shar
- F. Create a backup copy to the shared folde

- G. Sign in to the AWS Management Console and create an AMI from the backup copy Launch an EC2 instance that is based on the AMI.
- H. Create a managed-instance activation for a hybrid environment in AWS Systems Manager
- I. Download and install Systems Manager Agent on the on-premises VM Register the VM with Systems Manager to be a managed instance Use AWS Backup to create a snapshot of the VM and create an AMI
- J. Launch an EC2 instance that is based on the AMI

Answer: B

Explanation:

<https://docs.aws.amazon.com/vm-import/latest/userguide/vmimport-image-import.html>

- Export an OVF Template
- Create / use an Amazon S3 bucket for storing the exported images. The bucket must be in the Region where you want to import your VMs.
- Create an IAM role named vmimport.
- You'll use AWS CLI to run the import commands. <https://aws.amazon.com/premiumsupport/knowledge-center/import-instances/>

NEW QUESTION 132

- (Exam Topic 1)

A solutions architect is building a web application that uses an Amazon RDS for PostgreSQL DB instance The DB instance is expected to receive many more reads than writes The solutions architect needs to ensure that the large amount of read traffic can be accommodated and that the DB instance is highly available. Which steps should the solutions architect take to meet these requirements? (Select THREE.)

- A. Create multiple read replicas and put them into an Auto Scaling group
- B. Create multiple read replicas in different Availability Zones.
- C. Create an Amazon Route 53 hosted zone and a record set for each read replica with a TTL and a weighted routing policy
- D. Create an Application Load Balancer (ALB) and put the read replicas behind the ALB.
- E. Configure an Amazon CloudWatch alarm to detect a failed read replica Set the alarm to directly invoke an AWS Lambda function to delete its Route 53 record set.
- F. Configure an Amazon Route 53 health check for each read replica using its endpoint

Answer: BCF

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/requests-rds-read-replicas/>

You can use Amazon Route 53 weighted record sets to distribute requests across your read replicas. Within a Route 53 hosted zone, create individual record sets for each DNS endpoint associated with your read replicas and give them the same weight. Then, direct requests to the endpoint of the record set. You can incorporate Route 53 health checks to be sure that Route 53 directs traffic away from unavailable read replicas

NEW QUESTION 137

- (Exam Topic 1)

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 configuration
- 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 configuration
- G. Host the application and presentation 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 nodes
- 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

Explanation:

Aurora would improve availability that can replicate to multiple AZ (6 copies). Auto scaling would improve the performance together with a ALB. AppStream is like Citrix that deliver hosted Apps to users.

NEW QUESTION 139

- (Exam Topic 1)

A company wants to migrate its corporate data center from on premises to the AWS Cloud. The data center includes physical servers and VMs that use VMware and Hyper-V. An administrator needs to select the correct services to collect data (or the initial migration discovery process. The data format should be supported by AWS Migration Hub. The company also needs the ability to generate reports from the data.

Which solution meets these requirements?

- A. Use the AWS Agentless Discovery Connector for data collection on physical servers and all VM
- B. Store the collected data in Amazon S3. Query the data with S3 Select
- C. Generate reports by using Kibana hosted on Amazon EC2.
- D. Use the AWS Application Discovery Service agent for data collection on physical servers and all VMs. Store the collected data in Amazon Elastic File System (Amazon EFS). Query the data and generate reports with Amazon Athena.
- E. Use the AWS Application Discovery Service agent for data collection on physical servers and Hyper-V
- F. Use the AWS Agentless Discovery Connector for data collection on VMware
- G. Store the collected data in Amazon S3. Query the data with Amazon Athena

- H. Generate reports by using Amazon QuickSight.
- I. Use the AWS Systems Manager agent for data collection on physical server
- J. Use the AWS Agentless Discovery Connector for data collection on all VM
- K. Store, query, and generate reports from the collected data by using Amazon Redshift.

Answer: C

Explanation:

<https://docs.aws.amazon.com/application-discovery/latest/userguide/discovery-agent.html> <https://docs.aws.amazon.com/application-discovery/latest/userguide/discovery-connector.html>

NEW QUESTION 141

- (Exam Topic 1)

A company has a data lake in Amazon S3 that needs to be accessed by hundreds of applications across many AWS accounts. The company's information security policy states that the S3 bucket must not be accessed over the public internet and that each application should have the minimum permissions necessary to function.

To meet these requirements, a solutions architect plans to use an S3 access point that is restricted to specific VPCs for each application.

Which combination of steps should the solutions architect take to implement this solution? (Select TWO.)

- A. Create an S3 access point for each application in the AWS account that owns the S3 bucket
- B. Configure each access point to be accessible only from the application's VPC
- C. Update the bucket policy to require access from an access point.
- D. Create an interface endpoint for Amazon S3 in each application's VPC
- E. Configure the endpoint policy to allow access to an S3 access point
- F. Create a VPC gateway attachment for the S3 endpoint.
- G. Create a gateway endpoint for Amazon S3 in each application's VPC
- H. Configure the endpoint policy to allow access to an S3 access point
- I. Specify the route table that is used to access the access point.
- J. Create an S3 access point for each application in each AWS account and attach the access points to the S3 bucket
- K. Configure each access point to be accessible only from the application's VPC
- L. Update the bucket policy to require access from an access point.
- M. Create a gateway endpoint for Amazon S3 in the data lake's VPC
- N. Attach an endpoint policy to allow access to the S3 bucket
- O. Specify the route table that is used to access the bucket.

Answer: AC

Explanation:

<https://joe.blog.freemansoft.com/2020/04/protect-data-in-cloud-with-s3-access.html> <https://aws.amazon.com/s3/features/access-points/>
<https://aws.amazon.com/s3/features/access-points/>
&
<https://aws.amazon.com/blogs/storage/managing-amazon-s3-access-with-vpc-endpoints-and-s3-access-points/>

NEW QUESTION 143

- (Exam Topic 1)

A company's AWS architecture currently uses access keys and secret access keys stored on each instance to access AWS services. Database credentials are hard-coded on each instance. SSH keys for command-line remote access are stored in a secured Amazon S3 bucket. The company has asked its solutions architect to improve the security posture of the architecture without adding operational complexity.

Which combination of steps should the solutions architect take to accomplish this? (Select THREE.)

- A. Use Amazon EC2 instance profiles with an IAM role.
- B. Use AWS Secrets Manager to store access keys and secret access keys.
- C. Use AWS Systems Manager Parameter Store to store database credentials.
- D. Use a secure fleet of Amazon EC2 bastion hosts (or remote access).
- E. Use AWS KMS to store database credentials.
- F. Use AWS Systems Manager Session Manager for remote access

Answer: ACF

Explanation:

<https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager.html>

NEW QUESTION 147

- (Exam Topic 1)

A company has implemented an ordering system using an event-driven architecture. During initial testing, the system stopped processing orders. Further analysis revealed that one order message in an Amazon Simple Queue Service (Amazon SQS) standard queue was causing an error on the backend and blocking all subsequent order messages. The visibility timeout of the queue is set to 30 seconds, and the backend processing timeout is set to 10 seconds. A solutions architect needs to analyze faulty order messages and ensure that the system continues to process subsequent messages.

Which step should the solutions architect take to meet these requirements?

- A. Increase the backend processing timeout to 30 seconds to match the visibility timeout
- B. Reduce the visibility timeout of the queue to automatically remove the faulty message
- C. Configure a new SQS FIFO queue as a dead-letter queue to isolate the faulty messages
- D. Configure a new SQS standard queue as a dead-letter queue to isolate the faulty messages.

Answer: D

NEW QUESTION 150

- (Exam Topic 1)

A company wants to control its cost of Amazon Athena usage. The company has allocated a specific monthly budget for Athena usage. A solutions architect must design a solution that will prevent the company from exceeding the budgeted amount. Which solution will meet these requirements?

- A. Use AWS Budget
- B. Create an alarm (or when the cost of Athena usage reaches the budgeted amount for the month)
- C. Configure AWS Budgets actions to deactivate Athena until the end of the month.
- D. Use Cost Explorer to create an alert for when the cost of Athena usage reaches the budgeted amount for the month
- E. Configure Cost Explorer to publish notifications to an Amazon Simple Notification Service (Amazon SNS) topic.
- F. Use AWS Trusted Advisor to track the cost of Athena usage
- G. Configure an Amazon EventBridge (Amazon CloudWatch Events) rule to deactivate Athena until the end of the month whenever the cost reaches the budgeted amount for the month
- H. Use Athena workgroups to set a limit on the amount of data that can be scanned
- I. Set a limit that is appropriate for the monthly budget and the current pricing for Athena.

Answer: D

NEW QUESTION 152

- (Exam Topic 1)

A company runs a popular web application in an on-premises data center. The application receives four million views weekly. The company expects traffic to increase by 200% because of an advertisement that will be published soon.

The company needs to decrease the load on the origin before the increase of traffic occurs. The company does not have enough time to move the entire application to the AWS Cloud.

Which solution will meet these requirements?

- A. Create an Amazon CloudFront content delivery network (CDN). Enable query forwarding to the origin. Create a managed cache policy that includes query string
- B. Use an on-premises load balancer as the origin
- C. Offload the DNS querying to AWS to handle CloudFront CDN traffic.
- D. Create an Amazon CloudFront content delivery network (CDN) that uses a Real Time Messaging Protocol (RTMP) distribution
- E. Enable query forwarding to the origin
- F. Use an on-premises load balancer as the origin
- G. Offload the DNS querying to AWS to handle CloudFront CDN traffic.
- H. Create an accelerator in AWS Global Accelerator
- I. Add listeners for HTTP and HTTPS TCP ports. Create an endpoint group
- J. Create a Network Load Balancer (NLB), and attach it to the endpoint group
- K. Point the NLB to the on-premises server
- L. Offload the DNS querying to AWS to handle AWS Global Accelerator traffic.
- M. Create an accelerator in AWS Global Accelerator
- N. Add listeners for HTTP and HTTPS TCP ports. Create an endpoint group
- O. Create an Application Load Balancer (ALB), and attach it to the endpoint group
- P. Point the ALB to the on-premises server
- Q. Offload the DNS querying to AWS to handle AWS Global Accelerator traffic.

Answer: D

NEW QUESTION 157

- (Exam Topic 1)

A multimedia company needs to deliver its video-on-demand (VOD) content to its subscribers in a cost-effective way. The video files range in size from 1-15 GB and are typically viewed frequently for the first 6 months after creation, and then access decreases considerably. The company requires all video files to remain immediately available for subscribers. There are now roughly 30,000 files, and the company anticipates doubling that number over time.

What is the MOST cost-effective solution for delivering the company's VOD content?

- A. Store the video files in an Amazon S3 bucket using S3 Intelligent-Tiering
- B. Use Amazon CloudFront to deliver the content with the S3 bucket as the origin.
- C. Use AWS Elemental MediaConvert and store the adaptive bitrate video files in Amazon S3. Configure an AWS Elemental MediaPackage endpoint to deliver the content from Amazon S3.
- D. Store the video files in Amazon Elastic File System (Amazon EFS) Standard
- E. Enable EFS lifecycle management to move the video files to EFS Infrequent Access after 6 months
- F. Create an Amazon EC2 Auto Scaling group behind an Elastic Load Balancer to deliver the content from Amazon EFS.
- G. Store the video files in Amazon S3 Standard
- H. Create S3 Lifecycle rules to move the video files to S3 Standard-Infrequent Access (S3 Standard-IA) after 6 months and to S3 Glacier Deep Archive after 1 year
- I. Use Amazon CloudFront to deliver the content with the S3 bucket as the origin.

Answer: A

Explanation:

<https://d1.awsstatic.com/whitepapers/amazon-cloudfront-for-media.pdf> <https://aws.amazon.com/solutions/implementations/video-on-demand-on-aws/>

NEW QUESTION 162

- (Exam Topic 1)

A team collects and routes behavioral data for an entire company. The company runs a Multi-AZ VPC environment with public subnets, private subnets, and an internet gateway. Each public subnet also contains a NAT gateway. Most of the company's applications read from and write to Amazon Kinesis Data Streams. Most of the workloads are in private subnets.

A solutions architect must review the infrastructure. The solutions architect needs to reduce costs and maintain the function of the applications. The solutions architect uses Cost Explorer and notices that the cost in the EC2-Other category is consistently high. A further review shows that NatGateway-Bytes charges are increasing the cost in the EC2-Other category.

What should the solutions architect do to meet these requirements?

- A. Enable VPC Flow Log

- B. Use Amazon Athena to analyze the logs for traffic that can be remove
- C. Ensure that security groups are Mocking traffic that is responsible for high costs.
- D. Add an interface VPC endpoint for Kinesis Data Streams to the VP
- E. Ensure that applications have the correct IAM permissions to use the interface VPC endpoint.
- F. Enable VPC Flow Logs and Amazon Detective Review Detective findings for traffic that is not related to Kinesis Data Streams Configure security groups to block that traffic
- G. Add an interface VPC endpoint for Kinesis Data Streams to the VP
- H. Ensure that the VPC endpoint policy allows traffic from the applications.

Answer: D

Explanation:

<https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-endpoints-access.html>

<https://aws.amazon.com/premiumsupport/knowledge-center/vpc-reduce-nat-gateway-transfer-costs/>

VPC endpoint policies enable you to control access by either attaching a policy to a VPC endpoint or by using additional fields in a policy that is attached to an IAM user, group, or role to restrict access to only occur via the specified VPC endpoint

NEW QUESTION 164

- (Exam Topic 1)

A public retail web application uses an Application Load Balancer (ALB) in front of Amazon EC2 instances running across multiple Availability Zones (AZs) in a Region backed by an Amazon RDS MySQL Multi-AZ deployment. Target group health checks are configured to use HTTP and pointed at the product catalogue page. Auto Scaling is configured to maintain the web fleet size based on the ALB health check.

Recently, the application experienced an outage. Auto Scaling continuously replaced the instances during the outage. A subsequent investigation determined that the web server metrics were within the normal range, but the database tier was experiencing high load, resulting in severely elevated query response times.

Which of the following changes together would remediate these issues while improving monitoring capabilities for the availability and functionality of the entire application stack for future growth? (Select TWO.)

- A. Configure read replicas for Amazon RDS MySQL and use the single reader endpoint in the web application to reduce the load on the backend database tier.
- B. Configure the target group health check to point at a simple HTML page instead of a product catalog page and the Amazon Route 53 health check against the product page to evaluate full application functionalit
- C. Configure Amazon CloudWatch alarms to notify administrators when the site fails.
- D. Configure the target group health check to use a TCP check of the Amazon EC2 web server and theAmazon Route 53 health check against the product page to evaluate full application functionalit
- E. Configure Amazon CloudWatch alarms to notify administrators when the site fails.
- F. Configure an Amazon CloudWatch alarm for Amazon RDS with an action to recover a high-load, impaired RDS instance in the database tier.
- G. Configure an Amazon ElastiCache cluster and place it between the web application and RDS MySQL instances to reduce the load on the backend database tier.

Answer: BE

Explanation:

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/health-checks-types.html>

NEW QUESTION 165

- (Exam Topic 1)

A company has an internal application running on AWS that is used to track and process shipments in the company's warehouse. Currently, after the system receives an order, it emails the staff the information needed to ship a package. Once the package is shipped, the staff replies to the email and the order is marked as shipped.

The company wants to stop using email in the application and move to a serverless application model. Which architecture solution meets these requirements?

- A. Use AWS Batch to configure the different tasks required lo ship a packag
- B. Have AWS Batch trigger an AWS Lambda function that creates and prints a shipping labe
- C. Once that label is scanne
- D. as it leaves the warehouse, have another Lambda function move the process to the next step in the AWS Batch job.B.
- E. When a new order is created, store the order information in Amazon SQ
- F. Have AWS Lambda check the queue every 5 minutes and process any needed wor
- G. When an order needs to be shipped, have Lambda print the label in the warehous
- H. Once the label has been scanned, as it leaves the warehouse, have an Amazon EC2 instance update Amazon SOS.
- I. Update the application to store new order information in Amazon DynamoD
- J. When a new order is created, trigger an AWS Step Functions workflow, mark the orders as "in progress," and print a package label to the warehous
- K. Once the label has been scanned and fulfilled, the application will trigger an AWS Lambda function that will mark the order as shipped and complete the workflow.
- L. Store new order information in Amazon EF
- M. Have instances pull the new information from the NFS and send that information to printers in the warehous
- N. Once the label has been scanned, as it leaves the warehouse, have Amazon API Gateway call the instances to remove the order information from Amazon EFS.

Answer: C

NEW QUESTION 170

- (Exam Topic 1)

A company is migrating its three-tier web application from on-premises to the AWS Cloud. The company has the following requirements for the migration process:

- Ingest machine images from the on-premises environment.
- Synchronize changes from the on-premises environment to the AWS environment until the production cutover.
- Minimize downtime when executing the production cutover.
- Migrate the virtual machines' root volumes and data volumes.

Which solution will satisfy these requirements with minimal operational overhead?

- A. Use AWS Server Migration Service (SMS) to create and launch a replication job for each tier of the applicatio
- B. Launch instances from the AMIs created by AWS SM

- C. After initial testing, perform a final replication and create new instances from the updated AMIs.
- D. Create an AWS CLI VM Import/Export script to migrate each virtual machine
- E. Schedule the script to run incrementally to maintain changes in the application
- F. Launch instances from the AMIs created by VM Import/Export
- G. Once testing is done, rerun the script to do a final import and launch the instances from the AMIs.
- H. Use AWS Server Migration Service (SMS) to upload the operating system volume
- I. Use the AWS CLI import-snapshots command for the data volume
- J. Launch instances from the AMIs created by AWS SMS and attach the data volumes to the instances
- K. After initial testing, perform a final replication, launch new instances from the replicated AMI
- L. and attach the data volumes to the instances.
- M. Use AWS Application Discovery Service and AWS Migration Hub to group the virtual machines as an application
- N. Use the AWS CLI VM Import/Export script to import the virtual machines as AMI
- O. Schedule the script to run incrementally to maintain changes in the application
- P. Launch instances from the AMI
- Q. After initial testing, perform a final virtual machine import and launch new instances from the AMIs.

Answer: A

Explanation:

SMS can handle migrating the data volumes:

<https://aws.amazon.com/about-aws/whats-new/2018/09/aws-server-migration-service-adds-support-for-migrating-data-volumes/>

NEW QUESTION 174

- (Exam Topic 1)

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

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

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

Answer: B

Explanation:

People are being confused by the term 'reserve capacity'. This is not the same as an on-demand capacity reservation. This article by AWS clearly states that by 'reserving capacity' you are reserving the instances and reducing your costs. See <https://aws.amazon.com/aws-cost-management/aws-cost-optimization/reserved-instances/>

NEW QUESTION 176

- (Exam Topic 1)

A company runs an application on AWS. An AWS Lambda function uses credentials to authenticate to an Amazon RDS MySQL DB instance. A security risk assessment identified that these credentials are not frequently rotated. Also, encryption at rest is not enabled for the DB instance. The security team requires that both of these issues be resolved.

Which strategy should a solutions architect recommend to remediate these security risks?

- A. Configure the Lambda function to store and retrieve the database credentials in AWS Secrets Manager and enable rotation of the credential
- B. Take a snapshot of the DB instance and encrypt a copy of that snapshot
- C. Replace the DB instance with a new DB instance that is based on the encrypted snapshot.
- D. Enable IAM DB authentication on the DB instance
- E. Grant the Lambda execution role access to the DB instance
- F. Modify the DB instance and enable encryption.
- G. Enable IAM DB authentication on the DB instance
- H. Grant the Lambda execution role access to the DB instance
- I. Create an encrypted read replica of the DB instance
- J. Promote the encrypted read replica to be the new primary node.
- K. Configure the Lambda function to store and retrieve the database credentials as encrypted AWS Systems Manager Parameter Store parameter
- L. Create another Lambda function to automatically rotate the credential
- M. Create an encrypted read replica of the DB instance
- N. Promote the encrypted read replica to be the new primary node.

Answer: A

Explanation:

Parameter store can store DB credentials as secure string but CANNOT rotate secrets, hence, go with A + Cannot enable encryption on existing MySQL RDS instance, must create a new encrypted one from unencrypted snapshot.

<https://aws.amazon.com/blogs/security/rotate-amazon-rds-database-credentials-automatically-with-aws-secrets-manager/> Encrypting a unencrypted instance of DB or creating an encrypted replica of an unencrypted DB instance are not possible. Hence A is the only solution possible.

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html#Overview.Encryption>.

NEW QUESTION 181

- (Exam Topic 1)

A finance company hosts a data lake in Amazon S3. The company receives financial data records over SFTP each night from several third parties. The company runs its own SFTP server on an Amazon EC2 instance in a public subnet of a VPC. After the files are uploaded, they are moved to the data lake by a cron job that runs on the same instance. The SFTP server is reachable on DNS sftp.exam.com through the use of Amazon Route 53.

What should a solutions architect do to improve the reliability and scalability of the SFTP solution?

- A. Move the EC2 instance into an Auto Scaling group

- B. Place the EC2 instance behind an Application Load Balancer (ALB). Update the DNS record sftp.example.com in Route 53 to point to the ALB.
- C. Migrate the SFTP server to AWS Transfer for SFT
- D. Update the DNS record sftp.example.com in Route 53 to point to the server endpoint hostname.
- E. Migrate the SFTP server to a file gateway in AWS Storage Gatewa
- F. Update the DNS record sflp.example.com in Route 53 to point to the file gateway endpoint.
- G. Place the EC2 instance behind a Network Load Balancer (NLB). Update the DNS record sftp.example.com in Route 53 to point to the NLB.

Answer: B

NEW QUESTION 182

- (Exam Topic 1)

An ecommerce website running on AWS uses an Amazon RDS for MySQL DB instance with General Purpose SSD storage. The developers chose an appropriate instance type based on demand, and configured 100 GB of storage with a sufficient amount of free space.

The website was running smoothly for a few weeks until a marketing campaign launched. On the second day of the campaign, users reported long wait times and time outs. Amazon CloudWatch metrics indicated that both reads and writes to the DB instance were experiencing long response times. The CloudWatch metrics show 40% to 50% CPU and memory utilization, and sufficient free storage space is still available. The application server logs show no evidence of database connectivity issues.

What could be the root cause of the issue with the marketing campaign?

- A. It exhausted the I/O credit balance due to provisioning low disk storage during the setup phase.
- B. It caused the data in the tables to change frequently, requiring indexes to be rebuilt to optimize queries.
- C. It exhausted the maximum number of allowed connections to the database instance.
- D. It exhausted the network bandwidth available to the RDS for MySQL DB instance.

Answer: A

Explanation:

"When using General Purpose SSD storage, your DB instance receives an initial I/O credit balance of 5.4 million I/O credits. This initial credit balance is enough to sustain a burst performance of 3,000 IOPS for 30 minutes."

<https://aws.amazon.com/blogs/database/how-to-use-cloudwatch-metrics-to-decide-between-general-purpose-or>

NEW QUESTION 183

- (Exam Topic 1)

A solutions architect at a large company needs to set up network security for outbound traffic to the internet from all AWS accounts within an organization m AWS Organizations The organization has more than 100 AWS accounts, and the accounts route to each other by using a centralized AWS Transit Gateway. Each account has both an internet gateway and a NAT gateway for outbound traffic to the interne) The company deploys resources only Into a single AWS Region The company needs the ability to add centrally managed rule-based filtering on all outbound traffic to the internet for all AWS accounts in the organization The peak load of outbound traffic will not exceed 25 Gbps in each Availability Zone

Which solution meets these requirements?

- A. Creates a new VPC for outbound traffic to the internet Connect the existing transit gateway to the new VPC Configure a new NAT gateway Create an Auto Scaling group of Amazon EC2 Instances that run an open-source internet proxy for rule-based filtering across all Availability Zones in the Region Modify all default routes to point to the proxy's Auto Scaling group
- B. Create a new VPC for outbound traffic to the internet Connect the existing transit gateway to the new VPC Configure a new NAT gateway Use an AWS Network Firewall firewall for rule-based filtering Create Network Firewall endpoints In each Availability Zone Modify all default routes to point to the Network Firewall endpoints
- C. Create an AWS Network Firewall firewal for rule-based filtering in each AWS account Modify all default routes to point to the Network Firewall firewalls in each account.
- D. In each AWS account, create an Auto Scaling group of network-optimized Amazon EC2 instances that run an open-source internet proxy for rule-based filtering Modify all default routes to point to the proxy's Auto Scaling group.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/networking-and-content-delivery/deployment-models-for-aws-network-firewall/>

<https://aws.amazon.com/blogs/networking-and-content-delivery/deploy-centralized-traffic-filtering-using-aws-n>

NEW QUESTION 186

.....

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-C02 Exam with Our Prep Materials Via below:

<https://www.certleader.com/SAP-C02-dumps.html>