

Exam Questions SAP-C02

AWS Certified Solutions Architect - Professional

<https://www.2passeasy.com/dumps/SAP-C02/>



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 on several Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer. The load on the application varies throughout the day, and EC2 instances are scaled in and out on a regular basis. Log files from the EC2 instances are copied to a central Amazon S3 bucket every 15 minutes. The security team discovers that log files are missing from some of the terminated EC2 instances.

Which set of actions will ensure that log files are copied to the central S3 bucket from the terminated EC2 instances?

- A. Create a script to copy log files to Amazon S3, and store the script in a file on the EC2 instance
- B. Create an Auto Scaling lifecycle hook and an Amazon EventBridge (Amazon CloudWatch Events) rule to detect lifecycle events from the Auto Scaling group
- C. Invoke an AWS Lambda function on the autoscaling:EC2_INSTANCE_TERMINATING transition to send ABANDON to the Auto Scaling group to prevent termination, run the script to copy the log files, and terminate the instance using the AWS SDK.
- D. Create an AWS Systems Manager document with a script to copy log files to Amazon S3. Create an Auto Scaling lifecycle hook and an Amazon EventBridge (Amazon CloudWatch Events) rule to detect lifecycle events from the Auto Scaling group
- E. Invoke an AWS Lambda function on the autoscaling:EC2_INSTANCE_TERMINATING transition to call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send CONTINUE to the Auto Scaling group to terminate the instance.
- F. Change the log delivery rate to every 5 minutes
- G. Create a script to copy log files to Amazon S3, and add the script to EC2 instance user data
- H. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to detect EC2 instance termination
- I. Invoke an AWS Lambda function from the EventBridge (CloudWatch Events) rule that uses the AWS CLI to run the user-data script to copy the log files and terminate the instance.
- J. Create an AWS Systems Manager document with a script to copy log files to Amazon S3. Create an Auto Scaling lifecycle hook that publishes a message to an Amazon Simple Notification Service (Amazon SNS) topic
- K. From the SNS notification, call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send ABANDON to the Auto Scaling group to terminate the instance.

Answer: B

Explanation:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/adding-lifecycle-hooks.html>

- Refer to Default Result section - If the instance is terminating, both abandon and continue allow the instance to terminate. However, abandon stops any remaining actions, such as other lifecycle hooks, and continue allows any other lifecycle hooks to complete.

<https://aws.amazon.com/blogs/infrastructure-and-automation/run-code-before-terminating-an-ec2-auto-scaling-i> <https://github.com/aws-samples/aws-lambda-lifecycle-hooks-function>

<https://github.com/aws-samples/aws-lambda-lifecycle-hooks-function/blob/master/cloudformation/template.yam>

NEW QUESTION 7

- (Exam Topic 1)

A company needs to architect a hybrid DNS solution. This solution will use an Amazon Route 53 private hosted zone for the domain cloud.example.com for the resources stored within VPCs.

The company has the following DNS resolution requirements:

- On-premises systems should be able to resolve and connect to cloud.example.com.
- All VPCs should be able to resolve cloud.example.com.

There is already an AWS Direct Connect connection between the on-premises corporate network and AWS Transit Gateway. Which architecture should the company use to meet these requirements with the HIGHEST performance?

- A. Associate the private hosted zone to all the VPC
- B. Create a Route 53 inbound resolver in the shared services VP
- C. Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver.
- D. Associate the private hosted zone to all the VPC
- E. Deploy an Amazon EC2 conditional forwarder in the shared services VP
- F. Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the conditional forwarder.
- G. Associate the private hosted zone to the shared services VP
- H. Create a Route 53 outbound resolver in the shared services VP
- I. Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the outbound resolver.
- J. Associate the private hosted zone to the shared services VP
- K. Create a Route 53 inbound resolver in the shared services VP
- L. Attach the shared services VPC to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver.

Answer: D

Explanation:

<https://aws.amazon.com/blogs/networking-and-content-delivery/centralized-dns-management-of-hybrid-cloud-w>

NEW QUESTION 8

- (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
- I. application, and database tiers in both regions with equivalent capacity
- J. Activate the primary database in one region only and the standby database in the other region
- K. 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 9

- (Exam Topic 1)

A company wants to move a web application to AWS. The application stores session information locally on each web server, which will make auto scaling difficult. As part of the migration, the application will be rewritten to decouple the session data from the web servers. The company requires low latency, scalability, and availability.

Which service will meet the requirements for storing the session information in the MOST cost-effective way?

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

Answer: D

Explanation:

<https://aws.amazon.com/caching/session-management/>

Building real-time apps across versatile use cases like gaming, geospatial service, caching, session stores, or queuing, with advanced data structures, replication, and point-in-time snapshot support. Memcached: Building a simple, scalable caching layer for your data-intensive apps. <https://aws.amazon.com/elasticache/>

NEW QUESTION 10

- (Exam Topic 1)

A company has many AWS accounts and uses AWS Organizations to manage all of them. A solutions architect must implement a solution that the company can use to share a common network across multiple accounts.

The company's infrastructure team has a dedicated infrastructure account that has a VPC. The infrastructure team must use this account to manage the network. Individual accounts cannot have the ability to manage their own networks. However, individual accounts must be able to create AWS resources within subnets. Which combination of actions should the solutions architect perform to meet these requirements? (Select TWO.)

- A. Create a transit gateway in the infrastructure account.
- B. Enable resource sharing from the AWS Organizations management account.
- C. Create VPCs in each AWS account within the organization in AWS Organization
- D. Configure the VPCs to share the same CIDR range and subnets as the VPC in the infrastructure account
- E. Peer the VPCs in each individual account with the VPC in the infrastructure account,
- F. Create a resource share in AWS Resource Access Manager in the infrastructure account
- G. Select the specific AWS Organizations OU that will use the shared network
- H. Select each subnet to associate with the resource share.
- I. Create a resource share in AWS Resource Access Manager in the infrastructure account
- J. Select the specific AWS Organizations OU that will use the shared network
- K. Select each prefix list to associate with the resource share.

Answer: CE

Explanation:

<https://docs.aws.amazon.com/vpc/latest/userguide/sharing-managed-prefix-lists.html>

NEW QUESTION 10

- (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 Region
- 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 Region
- 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 overhead and automatic scaling

NEW QUESTION 15

- (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 multivalued 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 Region
- 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

O. 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 17

- (Exam Topic 1)

A solutions architect has an operational workload deployed on Amazon EC2 instances in an Auto Scaling group. The VPC architecture spans two Availability Zones (AZ) with a subnet in each that the Auto Scaling group is targeting. The VPC is connected to an on-premises environment and connectivity cannot be interrupted. The maximum size of the Auto Scaling group is 20 instances in service. The VPC IPv4 addressing is as follows:

VPC CIDR: 10.0.0.0/23

AZ1 subnet CIDR: 10.0.0.0/24 AZ2 subnet CIDR: 10.0.1.0/24

Since deployment, a third AZ has become available in the Region. The solutions architect wants to adopt the new AZ without adding additional IPv4 address space and without service downtime.

Which solution will meet these requirements?

- A. Update the Auto Scaling group to use the AZ2 subnet only
- B. Delete and re-create the AZ1 subnet using half the previous address space
- C. Adjust the Auto Scaling group to also use the new AZ1 subnet
- D. When the instances are healthy, adjust the Auto Scaling group to use the AZ1 subnet only
- E. Remove the current AZ2 subnet
- F. Create a new AZ2 subnet using the second half of the address space from the original AZ1 subnet
- G. Create a new AZ3 subnet using half the original AZ2 subnet address space, then update the Auto Scaling group to target all three new subnets.
- H. Terminate the EC2 instances in the AZ1 subnet
- I. Delete and re-create the AZ1 subnet using half the address space
- J. Update the Auto Scaling group to use this new subnet
- K. Repeat this for the second AZ
- L. Define a new subnet in AZ3, then update the Auto Scaling group to target all three new subnets.
- M. Create a new VPC with the same IPv4 address space and define three subnets, with one for each AZ
- N. Update the existing Auto Scaling group to target the new subnets in the new VPC.
- O. Update the Auto Scaling group to use the AZ2 subnet only
- P. Update the AZ1 subnet to have half the previous address space
- Q. Adjust the Auto Scaling group to also use the AZ1 subnet again
- R. When the instances are healthy, adjust the Auto Scaling group to use the AZ1 subnet only
- S. Update the current AZ2 subnet and assign the second half of the address space from the original AZ1 subnet
- T. Create a new AZ3 subnet using half the original AZ2 subnet address space, then update the Auto Scaling group to target all three new subnets.

Answer: A

Explanation:

https://aws.amazon.com/premiumsupport/knowledge-center/vpc-ip-address-range/?nc1=h_ls

It's not possible to modify the IP address range of an existing virtual private cloud (VPC) or subnet. You must delete the VPC or subnet, and then create a new VPC or subnet with your preferred CIDR block.

NEW QUESTION 20

- (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. Use the gateway.
- E. Recreate the data on 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 22

- (Exam Topic 1)

A company that is developing a mobile game is making game assets available in two AWS Regions. Game assets are served from a set of Amazon EC2 instances behind an Application Load Balancer (ALB) in each Region. The company requires game assets to be fetched from the closest Region. If game assets become unavailable in the closest Region, they should be fetched from the other Region.

What should a solutions architect do to meet these requirements?

- A. Create an Amazon CloudFront distribution

- B. Create an origin group with one origin for each AL
- C. Set one of the origins as primary.
- D. Create an Amazon Route 53 health check for each AL
- E. Create a Route 53 failover routing record pointing to the two ALB
- F. Set the Evaluate Target Health value to Yes.
- G. Create two Amazon CloudFront distributions, each with one ALB as the origi
- H. Create an Amazon Route 53 failover routing record pointing to the two CloudFront distribution
- I. Set the Evaluate Target Health value to Yes.
- J. Create an Amazon Route 53 health check for each AL
- K. Create a Route 53 latency alias record pointing to the two ALB
- L. Set the Evaluate Target Health value to Yes.

Answer: D

Explanation:

Failover routing policy – Use when you want to configure active-passive failover. Latency routing policy – Use when you have resources in multiple AWS Regions and you want to route traffic to the region that provides the best latency. <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html>

NEW QUESTION 25

- (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 organizatio
- 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 resources Deploy AWS WAF rules by using an AWS CloudFormation stack set to target the same OUs where the AWS Config rule is applied.
- E. Create AWS WAF rules in the management account of the organization 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 accounts 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.
- F. Use AWS Control Tower to manage AWS WAF rules across accounts in the organization 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 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 29

- (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 datasharing without causing the institute that owns the bucket to assume all the costs for S3 requests and data transfers'?

- A. Ensure that all organizations in the partnership have AWS account
- B. In the account with the S3 bucket, create a cross-account role for each account in the partnership that allows read access to the dat
- C. Have the organizations assume and use that read role when accessing the data.
- D. Ensure that all organizations in the partnership have AWS account
- E. Create a bucket policy on the bucket that owns the data The policy should allow the accounts in the partnership read access to the bucke
- F. Enable Requester Pays on the bucke
- 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 dat
- M. Enable Requester Pays on the bucke
- 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 34

- (Exam Topic 1)

A travel company built a web application that uses Amazon Simple Email Service (Amazon SES) to send email notifications to users. The company needs to enable logging to help troubleshoot email delivery issues. The company also needs the ability to do searches that are based on recipient, subject, and time sent. Which combination of steps should a solutions architect take to meet these requirements? (Select TWO.)

- A. Create an Amazon SES configuration set with Amazon Kinesis Data Firehose as the destination
- B. Choose to send logs to an Amazon S3 bucket.
- C. Enable AWS CloudTrail logging
- D. Specify an Amazon S3 bucket as the destination for the logs.
- E. Use Amazon Athena to query the logs in the Amazon S3 bucket for recipient, subject, and time sent.
- F. Create an Amazon CloudWatch log group
- G. Configure Amazon SES to send logs to the log group
- H. Use Amazon Athena to query the logs in Amazon CloudWatch for recipient, subject, and time sent.

Answer: AC

Explanation:

<https://docs.aws.amazon.com/ses/latest/dg/event-publishing-retrieving-firehose.html>

To enable you to track your email sending at a granular level, you can set up Amazon SES to publish email sending events to Amazon CloudWatch, Amazon Kinesis Data Firehose, or Amazon Simple Notification Service based on characteristics that you define.

<https://docs.aws.amazon.com/ses/latest/dg/monitor-using-event-publishing.html>

<https://aws.amazon.com/getting-started/hands-on/build-serverless-real-time-data-processing-app-lambda-kinesis>

NEW QUESTION 38

- (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 42

- (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 detection
- D. Configure cross-Region snapshots of the DB instance to the DR Region every 5 minutes
- 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 Region
- 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 Region
- J. Use scheduled AWS Glue jobs to synchronize the DB instance with another DB instance in the DR Region
- 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 45

- (Exam Topic 1)

A solutions architect must analyze a company's Amazon EC2 instances and Amazon Elastic Block Store (Amazon EBS) volumes to determine whether the company is using resources efficiently. The company is

running several large, high-memory EC2 instances to host database clusters that are deployed in active/passive configurations. The utilization of these EC2 instances varies by the applications that use the databases, and the company has not identified a pattern.

The solutions architect must analyze the environment and take action based on the findings. Which solution meets these requirements MOST cost-effectively?

- A. Create a dashboard by using AWS Systems Manager OpsCenter. Configure visualizations for Amazon CloudWatch metrics that are associated with the EC2 instances and their EBS volumes. Review the dashboard periodically and identify usage patterns. Rightsize the EC2 instances based on the peaks in the metrics.
- B. Turn on Amazon CloudWatch detailed monitoring for the EC2 instances and their EBS volumes. Create and review a dashboard that is based on the metrics. Identify usage patterns. Rightsize the EC2 instances based on the peaks in the metrics.
- C. Install the Amazon CloudWatch agent on each of the EC2 instances. Turn on AWS Compute Optimizer, and let it run for at least 12 hours. Review the recommendations from Compute Optimizer, and rightsize the EC2 instances as directed.
- D. Sign up for the AWS Enterprise Support plan. Turn on AWS Trusted Advisor. Wait 12 hours. Review the recommendations from Trusted Advisor, and rightsize the EC2 instances as directed.

Answer: C

Explanation:

(<https://aws.amazon.com/compute-optimizer/pricing/> , <https://aws.amazon.com/systems-manager/pricing/>). <https://aws.amazon.com/compute-optimizer/>

NEW QUESTION 49

- (Exam Topic 1)

A company has 50 AWS accounts that are members of an organization in AWS Organizations. Each account contains multiple VPCs. The company wants to use AWS Transit Gateway to establish connectivity between the VPCs in each member account. Each time a new member account is created, the company wants to automate the process of creating a new VPC and a transit gateway attachment.

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

- A. From the management account, share the transit gateway with member accounts by using AWS Resource Access Manager
- B. From the management account, share the transit gateway with member accounts by using an AWS Organizations SCP
- C. Launch an AWS CloudFormation stack set from the management account that automatically creates a new VPC and a VPC transit gateway attachment in a member account
- D. Associate the attachment with the transit gateway in the management account by using the transit gateway ID.
- E. Launch an AWS CloudFormation stack set from the management account that automatically creates a new VPC and a peering transit gateway attachment in a member account
- F. Share the attachment with the transit gateway in the management account by using a transit gateway service-linked role.
- G. From the management account, share the transit gateway with member accounts by using AWS Service Catalog

Answer: AC

NEW QUESTION 53

- (Exam Topic 1)

A company has multiple AWS accounts as part of an organization created with AWS Organizations. Each account has a VPC in the us-east-2 Region and is used for either production or development workloads. Amazon EC2 instances across production accounts need to communicate with each other, and EC2 instances across development accounts need to communicate with each other, but production and development instances should not be able to communicate with each other.

To facilitate connectivity, the company created a common network account. The company used AWS Transit Gateway to create a transit gateway in the us-east-2 Region in the network account and shared the transit gateway with the entire organization by using AWS Resource Access Manager. Network administrators then attached VPCs in each account to the transit gateway, after which the EC2 instances were able to communicate across accounts. However, production and development accounts were also able to communicate with one another.

Which set of steps should a solutions architect take to ensure production traffic and development traffic are completely isolated?

- A. Modify the security groups assigned to development EC2 instances to block traffic from production EC2 instances
- B. Modify the security groups assigned to production EC2 instances to block traffic from development EC2 instances.
- C. Create a tag on each VPC attachment with a value of either production or development, according to the type of account being attached
- D. Using the Network Manager feature of AWS Transit Gateway, create policies that restrict traffic between VPCs based on the value of this tag.
- E. Create separate route tables for production and development traffic
- F. Delete each account's association and route propagation to the default AWS Transit Gateway route table
- G. Attach development VPCs to the development AWS Transit Gateway route table and production VPCs to the production route table, and enable automatic route propagation on each attachment.
- H. Create a tag on each VPC attachment with a value of either production or development, according to the type of account being attached
- I. Modify the AWS Transit Gateway routing table to route production tagged attachments to one another and development tagged attachments to one another.

Answer: C

Explanation:

<https://docs.aws.amazon.com/vpc/latest/tgw/vpc-tgw.pdf>

NEW QUESTION 54

- (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 RDS
- C. and creating several additional read replicas to handle the load during end of month
- D. Using Amazon CloudWatch with AWS Lambda to change the type, size, or IOPS of Amazon EBS volumes in the cluster based on a specific CloudWatch metric
- 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 55

- (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 58

- (Exam Topic 1)

An e-commerce company is revamping its IT infrastructure and is planning to use AWS services. The company's CIO has asked a solutions architect to design a simple, highly available, and loosely coupled order processing application. The application is responsible for receiving and processing orders before storing them in an Amazon DynamoDB table. The application has a sporadic traffic pattern and should be able to scale during marketing campaigns to process the orders with minimal delays.

Which of the following is the MOST reliable approach to meet the requirements?

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

Answer: B

Explanation:

Q: How does Amazon Kinesis Data Streams differ from Amazon SQS?

Amazon Kinesis Data Streams enables real-time processing of streaming big data. It provides ordering of records, as well as the ability to read and/or replay records in the same order to multiple Amazon Kinesis Applications. The Amazon Kinesis Client Library (KCL) delivers all records for a given partition key to the same record processor, making it easier to build multiple applications reading from the same Amazon Kinesis data stream (for example, to perform counting, aggregation, and filtering).

<https://aws.amazon.com/kinesis/data-streams/faqs/>

<https://aws.amazon.com/blogs/big-data/unite-real-time-and-batch-analytics-using-the-big-data-lambda-architect>

NEW QUESTION 59

- (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 62

- (Exam Topic 1)

A company standardized its method of deploying applications to AWS using AWS CodePipeline and AWS CloudFormation. The applications are in Typescript and Python. The company has recently acquired another business that deploys applications to AWS using Python scripts.

Developers from the newly acquired company are hesitant to move their applications under CloudFormation because it would require them to learn a new domain-specific language and eliminate their access to language features, such as looping.

How can the acquired applications quickly be brought up to deployment standards while addressing the developers' concerns?

- A. Create CloudFormation templates and re-use parts of the Python scripts as instance user data
- B. Use the AWS Cloud Development Kit (AWS CDK) to deploy the application using these templates
- C. Incorporate the AWS CDK into CodePipeline and deploy the application to AWS using these templates.
- D. Use a third-party resource provisioning engine inside AWS CodeBuild to standardize the deployment processes of the existing and acquired company
- E. Orchestrate the CodeBuild job using CodePipeline.
- F. Standardize on AWS OpsWorks

- G. Integrate OpsWorks with CodePipeline
- H. Have the developers create Chef recipes to deploy their applications on AWS.
- I. Define the AWS resources using Typescript or Python
- J. Use the AWS Cloud Development Kit (AWS CDK) to create CloudFormation templates from the developers' code, and use the AWS CDK to create CloudFormation stack
- K. Incorporate the AWS CDK as a CodeBuild job in CodePipeline.

Answer: D

NEW QUESTION 63

- (Exam Topic 1)

A company is using AWS Organizations to manage multiple accounts. Due to regulatory requirements, the company wants to restrict specific member accounts to certain AWS Regions, where they are permitted to deploy resources. The resources in the accounts must be tagged, enforced based on a group standard, and centrally managed with minimal configuration.

What should a solutions architect do to meet these requirements?

- A. Create an AWS Config rule in the specific member accounts to limit Regions and apply a tag policy.
- B. From the AWS Billing and Cost Management console, in the master account, disable Regions for the specific member accounts and apply a tag policy on the root.
- C. Associate the specific member accounts with the root
- D. Apply a tag policy and an SCP using conditions to limit Regions.
- E. Associate the specific member accounts with a new O
- F. Apply a tag policy and an SCP using conditions to limit Regions.

Answer: D

NEW QUESTION 64

- (Exam Topic 1)

A web application is hosted in a dedicated VPC that is connected to a company's on-premises data center over a Site-to-Site VPN connection. The application is accessible from the company network only. This is a temporary non-production application that is used during business hours. The workload is generally low with occasional surges.

The application has an Amazon Aurora MySQL provisioned database cluster on the backend. The VPC has an internet gateway and a NAT gateway attached. The web servers are in private subnets in an Auto Scaling group behind an Elastic Load Balancer. The web servers also upload data to an Amazon S3 bucket through the internet.

A solutions architect needs to reduce operational costs and simplify the architecture. Which strategy should the solutions architect use?

- A. Review the Auto Scaling group settings and ensure the scheduled actions are specified to operate the Amazon EC2 instances during business hours only
- B. Use 3-year scheduled Reserved Instances for the web server EC2 instance
- C. Detach the internet gateway and remove the NAT gateways from the VPC
- D. Use an Aurora Serverless database and set up a VPC endpoint for the S3 bucket.
- E. Review the Auto Scaling group settings and ensure the scheduled actions are specified to operate the Amazon EC2 instances during business hours only
- F. Detach the internet gateway and remove the NAT gateways from the VPC
- G. Use an Aurora Serverless database and set up a VPC endpoint for the S3 bucket, then update the network routing and security rules and policies related to the changes.
- H. Review the Auto Scaling group settings and ensure the scheduled actions are specified to operate the Amazon EC2 instances during business hours only
- I. Detach the internet gateway from the VPC, and use an Aurora Serverless database
- J. Set up a VPC endpoint for the S3 bucket, then update the network routing and security rules and policies related to the changes.
- K. Use 3-year scheduled Reserved Instances for the web server Amazon EC2 instance
- L. Remove the NAT gateways from the VPC, and set up a VPC endpoint for the S3 bucket
- M. Use Amazon
- N. CloudWatch and AWS Lambda to stop and start the Aurora DB cluster so it operates during business hours only
- O. Update the network routing and security rules and policies related to the changes.

Answer: B

Explanation:

The application is accessible from the company network only remove NAT and IGW, application - S3 with VPC endpoint. Non-Production application no need to go for Reserved instances

To build site-to-site vpn, you don't need internet gateway. Instead, customer gateway is needed.

<https://docs.aws.amazon.com/vpn/latest/s2svpn/SetUpVPNConnections.html#vpn-create-cgw>

NEW QUESTION 68

- (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 SQS queue to process the records, writing the results to a temporary location in Amazon S3. Trigger a final Lambda function once the SQS 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 SQS queue
- D. Configure an AWS Fargate container application to
- E. automatically scale to a single instance when the SQS 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 requirement

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 requirement

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 70

- (Exam Topic 1)

A company is running a tone-of-business (LOB) application on AWS to support its users. The application runs in one VPC, with a backup copy in a second VPC in a different AWS Region for disaster recovery. The company has a single AWS Direct Connect connection between its on-premises network and AWS. The connection terminates at a Direct Connect gateway.

All access to the application must originate from the company's on-premises network, and traffic must be encrypted in transit through the use of IPsec. The company is routing traffic through a VPN tunnel over the Direct Connect connection to provide the required encryption.

A business continuity audit determines that the Direct Connect connection represents a potential single point of failure for access to the application. The company needs to remediate this issue as quickly as possible.

Which approach will meet these requirements?

A. Order a second Direct Connect connection to a different Direct Connect location.

B. Terminate the second Direct Connect connection at the same Direct Connect gateway.

C. Configure an AWS Site-to-Site VPN connection over the internet. Terminate the VPN connection at a virtual private gateway in the secondary Region.

D. Create a transit gateway. Attach the VPCs to the transit gateway, and connect the transit gateway to the Direct Connect gateway. Configure an AWS Site-to-Site VPN connection, and terminate it at the transit gateway.

E. Create a transit gateway.

F. Attach the VPCs to the transit gateway, and connect the transit gateway to the Direct Connect gateway.

G. Order a second Direct Connect connection, and terminate it at the transit gateway.

Answer: C

Explanation:

Create a transit gateway. Attach the VPCs to the transit gateway, and connect the transit gateway to the Direct Connect gateway. Configure an AWS Site-to-Site VPN connection, and terminate it at the transit gateway.

<https://aws.amazon.com/premiumsupport/knowledge-center/dx-configure-dx-and-vpn-failover-tgw/>

All access to the application must originate from the company's on-premises network and traffic must be encrypted in transit through the use of IPsec. = need to use VPN.

NEW QUESTION 74

- (Exam Topic 1)

An education company is running a web application used by college students around the world. The application runs in an Amazon Elastic Container Service (Amazon ECS) cluster in an Auto Scaling group behind an Application Load Balancer (ALB). A system administrator detects a weekly spike in the number of failed login attempts, which overwhelm the application's authentication service. All the failed login attempts originate from about 500 different IP addresses that change each week. A solutions architect must prevent the failed login attempts from overwhelming the authentication service.

Which solution meets these requirements with the MOST operational efficiency?

A. Use AWS Firewall Manager to create a security group and security group policy to deny access from the IP addresses.

B. Create an AWS WAF web ACL with a rate-based rule, and set the rule action to Block.

C. Connect the web ACL to the ALB.

D. Use AWS Firewall Manager to create a security group and security group policy to allow access only to specific CIDR ranges.

E. Create an AWS WAF web ACL with an IP set match rule, and set the rule action to Block.

F. Connect the web ACL to the ALB.

Answer: B

Explanation:

<https://docs.aws.amazon.com/waf/latest/developerguide/waf-rule-statement-type-rate-based.html>

The IP set match statement inspects the IP address of a web request against a set of IP addresses and address ranges. Use this to allow or block web requests based on the IP addresses that the requests originate from. By default, AWS WAF uses the IP address from the web request origin, but you can configure the rule to use an HTTP header like X-Forwarded-For instead.

<https://docs.aws.amazon.com/waf/latest/developerguide/waf-rule-statement-type-ipset-match.html>

<https://docs.aws.amazon.com/waf/latest/developerguide/waf-rule-statement-type-rate-based.html>

NEW QUESTION 76

- (Exam Topic 1)

A solutions architect is designing an application to accept timesheet entries from employees on their mobile devices. Timesheets will be submitted weekly, with most of the submissions occurring on Friday. The data must be stored in a format that allows payroll administrators to run monthly reports. The infrastructure must be highly available and scale to match the rate of incoming data and reporting requests.

Which combination of steps meets these requirements while minimizing operational overhead? (Select TWO.)

- A. Deploy the application to Amazon EC2 On-Demand Instances With load balancing across multiple Availability Zone
- B. Use scheduled Amazon EC2 Auto Scaling to add capacity before the high volume of submissions on Fridays.
- C. Deploy the application in a container using Amazon Elastic Container Service (Amazon ECS) with load balancing across multiple Availability Zone
- D. Use scheduled Service Auto Scaling to add capacity before the high volume of submissions on Fridays.
- E. Deploy the application front end to an Amazon S3 bucket served by Amazon CloudFront
- F. Deploy the application backend using Amazon API Gateway with an AWS Lambda proxy integration.
- G. Store the timesheet submission data in Amazon Redshift
- H. Use Amazon QuickSight to generate the reports using Amazon Redshift as the data source.
- I. Store the timesheet submission data in Amazon S3. Use Amazon Athena and Amazon QuickSight to generate the reports using Amazon S3 as the data source.

Answer: AE

NEW QUESTION 81

- (Exam Topic 1)

A company is developing and hosting several projects in the AWS Cloud. The projects are developed across multiple AWS accounts under the same organization in AWS Organizations. The company requires the cost for cloud infrastructure to be allocated to the owning project. The team responsible for all of the AWS accounts has discovered that several Amazon EC2 instances are lacking the Project tag used for cost allocation.

Which actions should a solutions architect take to resolve the problem and prevent it from happening in the future? (Select THREE.)

- A. Create an AWS Config rule in each account to find resources with missing tags.
- B. Create an SCP in the organization with a deny action for ec2:RunInstances if the Project tag is missing.
- C. Use Amazon Inspector in the organization to find resources with missing tags.
- D. Create an IAM policy in each account with a deny action for ec2:RunInstances if the Project tag is missing.
- E. Create an AWS Config aggregator for the organization to collect a list of EC2 instances with the missing Project tag.
- F. Use AWS Security Hub to aggregate a list of EC2 instances with the missing Project tag.

Answer: BDE

NEW QUESTION 85

- (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 89

- (Exam Topic 1)

A company is building an image service on the web that will allow users to upload and search random photos. At peak usage, up to 10,000 users worldwide will upload their images. The service will then overlay text on the uploaded images, which will then be published on the company website.

Which design should a solutions architect implement?

- A. Store the uploaded images in Amazon Elastic File System (Amazon EFS). Send application log information about each image to Amazon CloudWatch Log
- B. Create a fleet of Amazon EC2 instances that use CloudWatch Logs to determine which images need to be processed
- C. Place processed images in another directory in Amazon EFS
- D. Enable Amazon CloudFront and configure the origin to be the one of the EC2 instances in the fleet.
- E. Store the uploaded images in an Amazon S3 bucket and configure an S3 bucket event notification to send a message to Amazon Simple Notification Service (Amazon SNS). Create a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB) to pull messages from Amazon SNS to process the images and place them in Amazon Elastic File System (Amazon EFS). Use Amazon CloudWatch metrics for the SNS message volume to scale out EC2 instances
- F. Enable Amazon CloudFront and configure the origin to be the ALB in front of the EC2 instances.
- G. Store the uploaded images in an Amazon S3 bucket and configure an S3 bucket event notification to send a message to the Amazon Simple Queue Service (Amazon SQS) queue
- H. Create a fleet of Amazon EC2 instances to pull messages from the SQS queue to process the images and place them in another S3 bucket
- I. Use Amazon CloudWatch metrics for queue depth to scale out EC2 instances
- J. Enable Amazon CloudFront and configure the origin to be the S3 bucket that contains the processed images.
- K. Store the uploaded images on a shared Amazon Elastic Block Store (Amazon EBS) volume mounted to a fleet of Amazon EC2 Spot instances
- L. Create an Amazon DynamoDB table that contains information about each uploaded image and whether it has been processed
- M. Use an Amazon EventBridge (Amazon CloudWatch Events) rule to scale out EC2 instances
- N. Enable Amazon CloudFront and configure the origin to reference an Elastic Load Balancer in front of the fleet of EC2 instances.

Answer: C

NEW QUESTION 92

- (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 94

- (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 issue1?

- A. Suspend the Auto Scaling group's HealthCheck scaling proces
- 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 grou
- E. Use Session Manager to log in to an instance that is marked as unhealthy
- F. Suspend the Auto Scaling group's Terminate proces
- 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 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 95

- (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 MySQ
- 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 96

- (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 AW5 Transit Gatewa
- B. Attach the shared VPC and the authorized business unit VPCs to the transit gatewa
- 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 tabl
- 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 (he option to require endpoint acceptanc
- G. Create a VPC endpoint in each of the business unit VPCs using the service name of the endpoint servic
- H. Accept authorized endpoint requests from the endpoint service console.
- I. Create a VPC peering connection from each business unit VPC to lthe shared VP
- J. Accept the VPC peering connections from the shared VPC consol
- 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 theauthorized business unit VPC
- M. Establish a Sile-to-Site VPN connection from the business unit VPCs to the shared VP
- 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 100

- (Exam Topic 1)

A company is running an application distributed over several Amazon EC2 instances in an Auto Seating 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 sourc
- 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 105

- (Exam Topic 1)

A company maintains a restaurant review website. The website is a single-page application where files are stored in Amazon S3 and delivered using Amazon CloudFront. The company receives several fake postings every day that are manually removed.

The security team has identified that most of the fake posts are from bots with IP addresses that have a bad reputation within the same global region. The team needs to create a solution to help restrict the bots from accessing the website.

Which strategy should a solutions architect use?

- A. Use AWS Firewall Manager to control the CloudFront distribution security setting
- B. Create a geographical block rule and associate it with Firewall Manager.
- C. Associate an AWS WAF web ACL with the CloudFront distributio
- D. Select the managed Amazon IP reputation rule group for the web ACL with a deny action.
- E. Use AWS Firewall Manager to control the CloudFront distribution security setting
- F. Select the managed Amazon IP reputation rule group and associate it with Firewall Manager with a deny action.
- G. Associate an AWS WAF web ACL with the CloudFront distributio
- H. Create a rule group for the web ACL with a geographical match statement with a deny action.

Answer: B

Explanation:

IP reputation rule groups allow you to block requests based on their source. Choose one or more of these rule groups if you want to reduce your exposure to BOTS!!!! traffic or exploitation attempts

The Amazon IP reputation list rule group contains rules that are based on Amazon internal threat intelligence. This is useful if you would like to block IP addresses typically associated with bots or other threats. Inspects for a list of IP addresses that have been identified as bots by Amazon threat intelligence.

NEW QUESTION 108

- (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 DLMJ) 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 111

- (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 115

- (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 connection
- 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 gateway
- M. Delete the existing private virtual interface from the existing connection. Create the second Direct Connect connection
- N. Create a new private virtual interface on each connection, and connect both private virtual interfaces to the transit gateway
- 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 117

- (Exam Topic 1)

A company has a project that is launching Amazon EC2 instances that are larger than required. The project's account cannot be part of the company's organization in AWS Organizations due to policy restrictions to keep this activity outside of corporate IT. The company wants to allow only the launch of t3.small EC2 instances by developers in the project's account. These EC2 instances must be restricted to the us-east-2 Region.

What should a solutions architect do to meet these requirements?

- A. Create a new developer account
- B. Move all EC2 instances, users, and assets into us-east-2. Add the account to the company's organization in AWS Organization
- C. Enforce a tagging policy that denotes Region affinity.
- D. Create an SCP that denies the launch of all EC2 instances except t3.small EC2 instances in us-east-2. Attach the SCP to the project's account.
- E. Create and purchase a t3.small EC2 Reserved Instance for each developer in us-east-2. Assign each developer a specific EC2 instance with their name as the tag.
- F. Create an IAM policy that allows the launch of only t3.small EC2 instances in us-east-2. Attach the policy to the roles and groups that the developers use in the project's account.

Answer: D

NEW QUESTION 120

- (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 replic
- F. Set the alarm to directly invoke an AWS Lambda function to delete its Route 53 record set.
- G. 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 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 124

- (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 (CIFSJ 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 Manage
- 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 AM
- 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 128

- (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-
- F. Use the AWS Agentless Discovery Connector for data collection on VMwar
- 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 131

- (Exam Topic 1)

A company is using AWS CodePipeline for the CI/CO of an application to an Amazon EC2 Auto Scaling group. All AWS resources are defined in AWS CloudFormation templates. The application artifacts are stored in an Amazon S3 bucket and deployed to the Auto Scaling group using instance user data scripts. As the application has become more complex, recent resource changes in the Cloud Formation templates have caused unplanned downtime. How should a solutions architect improve the CI/CD pipeline to reduce the likelihood that changes in the templates will cause downtime?

- A. Adapt the deployment scripts to detect and report CloudFormation error conditions when performing deployment
- B. Write test plans for a testing team to execute in a non-production environment before approving the change for production.
- C. Implement automated testing using AWS CodeBuild in a test environment
- D. Use CloudFormation changesets to evaluate changes before deployment
- E. Use AWS CodeDeploy to leverage blue/green deployment patterns to allow evaluations and the ability to revert changes, if needed.
- F. Use plugins for the integrated development environment (IDE) to check the templates for errors, and use the AWS CLI to validate that the templates are correct
- G. Adapt the deployment code to check for error conditions and generate notifications on error
- H. Deploy to a test environment and execute a manual test plan before approving the change for production.
- I. Use AWS CodeDeploy and a blue/green deployment pattern with CloudFormation to replace the user data deployment script
- J. Have the operators log in to running instances and go through a manual test plan to verify the application is running as expected.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/devops/performing-bluegreen-deployments-with-aws-codedeploy-and-auto-scaling/> When one adopts go infrastructure as code, we need to test the infrastructure code as well via automated testing, and revert to original if things are not performing correctly.

NEW QUESTION 134

- (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 139

- (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 removed
- C. Ensure that security groups are blocking traffic that is responsible for high costs.
- D. Add an interface VPC endpoint for Kinesis Data Streams to the VPC
- 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 VPC
- 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 143

- (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 the Amazon 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 146

- (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 149

- (Exam Topic 1)

A company wants to migrate a 30 TB Oracle data warehouse from on premises to Amazon Redshift The company used the AWS Schema Conversion Tool (AWS SCT) to convert the schema of the existing data warehouse to an Amazon Redshift schema The company also used a migration assessment report to identify manual tasks to complete.

The company needs to migrate the data to the new Amazon Redshift cluster during an upcoming data freeze period of 2 weeks The only network connection between the on-premises data warehouse and AWS is a 50 Mops internet connection

Which migration strategy meets these requirements?

- A. Create an AWS Database Migration Service (AWS DMS) replication instanc
- B. Authorize the public IP address of the replication instance to reach the data warehouse through the corporate firewall Create a migration task to run at the beginning of the data freeze period.
- C. Install the AWS SCT extraction agents on the on-premises server
- D. Define the extract, upload, and copy tasks to send the data to an Amazon S3 bucke
- E. Copy the data into the Amazon Redshift cluste
- F. Run the tasks at the beginning of the data freeze period.
- G. install the AWS SCT extraction agents on the on-premises server
- H. Create a Site-to-Site VPN connection Create an AWS Database Migration Service (AWS DMS) replication instance that is the appropriate size Authorize the IP address of the replication instance to be able to access the on-premises data warehouse through the VPN connection
- I. Create a job in AWS Snowball Edge to import data into Amazon S3 Install AWS SCT extraction agents on the on-premises servers Define the local and AWS Database Migration Service (AWS DMS) tasks to send the data to the Snowball Edge device When the Snowball Edge device is returned to AWS and the data is available in Amazon S3, run the AWS DMS subtask to copy the data to Amazon Redshift.

Answer: D

Explanation:

AWS Database Migration Service (AWS DMS) can use Snowball Edge and Amazon S3 to migrate large databases more quickly than by other methods
https://docs.aws.amazon.com/dms/latest/userguide/CHAP_LargeDBs.html
https://www.calctool.org/CALC/prof/computing/transfer_time

NEW QUESTION 150

- (Exam Topic 1)

A company uses AWS Transit Gateway for a hub-and-spoke model to manage network traffic between many VPCs. The company is developing a new service that must be able to send data at 100 Gbps. The company needs a faster connection to other VPCs in the same AWS Region. Which solution will meet these requirements?

- A. Establish VPC peering between the necessary VPC
- B. Ensure that all route tables are updated as required.
- C. Attach an additional transit gateway to the VPC
- D. Update the route tables accordingly.
- E. Create AWS Site-to-Site VPN connections that use equal-cost multi-path (ECMP) routing between the necessary VPCs.
- F. Create an additional attachment from the necessary VPCs to the existing transit gateway.

Answer: D

NEW QUESTION 153

- (Exam Topic 1)

A company is migrating applications from on premises to the AWS Cloud. These applications power the company's internal web forms. These web forms collect data for specific events several times each quarter. The web forms use simple SQL statements to save the data to a local relational database. Data collection occurs for each event, and the on-premises servers are idle most of the time. The company needs to minimize the amount of idle infrastructure that supports the web forms. Which solution will meet these requirements?

- A. Use Amazon EC2 Image Builder to create AMIs for the legacy server
- B. Use the AMIs to provision EC2 instances to recreate the applications in the AWS Cloud
- C. Place an Application Load Balancer (ALB) in front of the EC2 instance
- D. Use Amazon Route 53 to point the DNS names of the web forms to the ALB.
- E. Create one Amazon DynamoDB table to store data for all the data input Use the application form name as the table key to distinguish data item
- F. Create an Amazon Kinesis data stream to receive the data input and store the input in DynamoDB
- G. Use Amazon Route 53 to point the DNS names of the web forms to the Kinesis data stream's endpoint.
- H. Create Docker images for each server of the legacy web form application
- I. Create an Amazon Elastic Container Service (Amazon ECS) cluster on AWS Fargate
- J. Place an Application Load Balancer in front of the ECS cluster
- K. Use Fargate task storage to store the web form data.
- L. Provision an Amazon Aurora Serverless cluster
- M. Build multiple schemas for each web form's data storage
- N. Use Amazon API Gateway and an AWS Lambda function to recreate the data input form
- O. Use Amazon Route 53 to point the DNS names of the web forms to their corresponding API Gateway endpoint.

Answer: D

Explanation:

Provision an Amazon Aurora Serverless cluster. Build multiple schemas for each web forms data storage. Use Amazon API Gateway and an AWS Lambda function to recreate the data input forms. Use Amazon Route 53 to point the DNS names of the web forms to their corresponding API Gateway endpoint.

NEW QUESTION 155

- (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 157

- (Exam Topic 2)

A company that develops consumer electronics with offices in Europe and Asia has 60 TB of software images stored on premises in Europe. The company wants to transfer the images to an Amazon S3 bucket in the ap-northeast-1 Region. New software images are created daily and must be encrypted in transit. The company needs a solution that does not require custom development to automatically transfer all existing and new software images to Amazon S3.

What is the next step in the transfer process?

- A. Deploy an AWS DataSync agent and configure a task to transfer the images to the S3 bucket.
- B. Configure Amazon Kinesis Data Firehose to transfer the images using S3 Transfer Acceleration.
- C. Use an AWS Snowball device to transfer the images with the S3 bucket as the target.
- D. Transfer the images over a Site-to-Site VPN connection using the S3 API with multipart upload.

Answer: A

NEW QUESTION 159

- (Exam Topic 2)

A company has developed a new billing application that will be released in two weeks. Developers are testing the application running on 10 EC2 instances managed by an Auto Scaling group in subnet 172.31.0.0/24 within VPC A with CIDR block 172.31.0.0/16. The developers noticed connection timeout errors in the application logs while connecting to an Oracle database running on an Amazon EC2 instance in the same region within VPC B with CIDR block 172.50.0.0/16. The IP of the database instance is hard-coded in the application instances.

Which recommendations should a solutions architect present to the developers to solve the problem in a secure way with minimal maintenance and overhead?

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

Answer: C

NEW QUESTION 161

- (Exam Topic 2)

A company wants to allow its marketing team to perform SQL queries on customer records to identify market segments. The data is spread across hundreds of files. The records must be encrypted in transit and at rest. The team manager must have the ability to manage users and groups but no team members should have access to services or resources not required for the SQL queries. Additionally, administrators need to audit the queries made and receive notifications when a query violates rules defined by the security team.

AWS Organizations has been used to create a new account and an AWS IAM user with administrator permissions for the team manager. Which design meets these requirements?

- A. Apply a service control policy (SCP) that allows access to IAM Amazon RD
- B. and AWS CloudTrail Load customer records in Amazon RDS MySQL and train users to run queries using the AWS CL
- C. Stream the query logs to Amazon CloudWatch Logs from the RDS database instance Use a subscription filter with AWS Lambda functions to audit and alarm on queries against personal data
- D. Apply a service control policy (SCP) that denies access to all services except IAM Amazon Athena Amazon S3 and AWS CloudTrail Store customer record files in Amazon S3 and tram users to run queries using the CLI via Athena Analyze CloudTrail events to audit and alarm on queries against personal data
- E. Apply a service control policy (SCP) that denies access to all services except IAM Amazon DynamoD
- F. and AWS CloudTrail Store customer records in DynamoDB and train users to run queries using the AWS CLI Enable DynamoDB streams to track the queries that are issued and use an AWS Lambda function for real-time monitoring and alerting
- G. Apply a service control policy (SCP) that allows access to IAM Amazon Athena; Amazon S3, and AWS CloudTrail Store customer records as files in Amazon S3 and train users to leverage the Amazon S3 Select feature and run queries using the AWS CLI Enable S3 object-level logging and analyze CloudTrail events to audit and alarm on queries against personal data

Answer: B

NEW QUESTION 163

- (Exam Topic 2)

A company is running a two-tier web-based application in an on-premises data center. The application layer consists of a single server running a stateful application. The application connects to a PostgreSQL database running on a separate server. The application's user base is expected to grow significantly, so the company is migrating the application and database to AWS. The solution will use Amazon Aurora PostgreSQL, Amazon EC2 Auto Scaling, and Elastic Load Balancing.

Which solution will provide a consistent user experience that will allow the application and database tiers to scale?

- A. Enable Aurora Auto Scaling for Aurora Replica
- B. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled
- C. Enable Aurora Auto Scaling for Aurora writer
- D. Use an Application Load Balancer with the round robin routing algorithm and sticky sessions enabled
- E. Aurora Auto Scaling for Aurora Replica
- F. Use an Application Load Balancer with the round robin routing algorithm and sticky sessions enabled.
- G. Aurora Auto Scaling for Aurora writer
- H. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.

Answer: C

NEW QUESTION 168

- (Exam Topic 2)

A company runs a proprietary stateless ETL application on an Amazon EC2 Linux instance. The application is a Linux binary, and the source code cannot be modified. The application is single-threaded, uses 2 GB of RAM, and is highly CPU intensive. The application is scheduled to run every 4 hours and runs for up to 20 minutes. A solutions architect wants to revise the architecture for the solution.

Which strategy should the solutions architect use?

- A. Use AWS Lambda to run the applicatio

- B. Use Amazon CloudWatch Logs to invoke the Lambda function every 4 hours
- C. Use AWS Batch to run the application Use an AWS Step Functions state machine to invoke the AWS Batch job every 4 hours
- D. Use AWS Fargate to run the application Use Amazon EventBridge (Amazon CloudWatch Events) to invoke the Fargate task every 4 hours
- E. Use Amazon EC2 Spot Instances to run the application Use AWS CodeDeploy to deploy and run the application every 4 hours.

Answer: C

NEW QUESTION 173

- (Exam Topic 2)

A company wants to migrate its website from an on-premises data center onto AWS. At the same time it wants to migrate the website to a containerized microservice-based architecture to improve the availability and cost efficiency. The company's security policy states that privileges and network permissions must be configured according to best practice, using least privilege.

A solutions architect must create a containerized architecture that meets the security requirements and has deployed the application to an Amazon ECS cluster. What steps are required after the deployment to meet the requirements? (Select TWO.)

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

Answer: BE

NEW QUESTION 175

- (Exam Topic 2)

A company is running a three-tier web application in an on-premises data center. The frontend is served by an Apache web server, the middle tier is a monolithic Java application, and the storage tier is a PostgreSQL database.

During a recent marketing promotion, customers could not place orders through the application because the application crashed. An analysis showed that all three tiers were overloaded. The application became unresponsive, and the database reached its capacity limit because of read operations. The company already has several similar promotions scheduled in the near future.

A solutions architect must develop a plan for migration to AWS to resolve these issues. The solution must maximize scalability and must minimize operational effort.

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

- A. Refactor the frontend so that static assets can be hosted on Amazon S3. Use Amazon CloudFront to serve the frontend to customer
- B. Connect the frontend to the Java application.
- C. Rehost the Apache web server of the frontend on Amazon EC2 instances that are in an Auto Scaling group
- D. Use a load balancer in front of the Auto Scaling group
- E. Use Amazon Elastic File System (Amazon EFS) to host the static assets that the Apache web server needs.
- F. Rehost the Java application in an AWS Elastic Beanstalk environment that includes auto scaling.
- G. Refactor the Java application
- H. Develop a Docker container to run the Java application
- I. Use AWS Fargate to host the container.
- J. Use AWS Database Migration Service (AWS DMS) to replatform the PostgreSQL database to an Amazon Aurora PostgreSQL database
- K. Use Aurora Auto Scaling for read replicas.
- L. Rehost the PostgreSQL database on an Amazon EC2 instance that has twice as much memory as the on-premises server.

Answer: BCF

NEW QUESTION 180

- (Exam Topic 2)

A fleet of Amazon ECS instances is used to poll an Amazon SQS queue and update items in an Amazon DynamoDB database. Items in the table are not being updated, and the SQS queue is filling up. Amazon CloudWatch Logs are showing consistent 400 errors when attempting to update the table. The provisioned write capacity units are appropriately configured, and no throttling is occurring.

What is the LIKELY cause of the failure*?

- A. The ECS service was deleted
- B. The ECS configuration does not contain an Auto Scaling group
- C. The ECS instance task execution IAM role was modified
- D. The ECS task role was modified

Answer: D

NEW QUESTION 183

- (Exam Topic 2)

A company has built a high performance computing (HPC) cluster in AWS for a tightly coupled workload that generates a large number of shared files stored in Amazon EFS. The cluster was performing well when the number of Amazon EC2 instances in the cluster was 100. However, when the company increased the cluster size to 1,000 EC2 instances, overall performance was well below expectations.

Which collection of design choices should a solutions architect make to achieve the maximum performance from the HPC cluster? (Select THREE.)

- A. Ensure the HPC cluster is launched within a single Availability Zone.
- B. Launch the EC2 instances and attach elastic network interfaces in multiples of four.
- C. Select EC2 instance types with an Elastic Fabric Adapter (EFA) enabled
- D. Ensure the cluster is launched across multiple Availability Zones.
- E. Replace Amazon EFS with multiple Amazon EBS volumes in a RAID array.
- F. Replace Amazon EFS with Amazon FSx for Lustre.

Answer: ACE

NEW QUESTION 185

- (Exam Topic 2)

A large company has many business units. Each business unit has multiple AWS accounts for different purposes. The CIO of the company sees that each business unit has data that would be useful to share with other parts of the company. In total, there are about 10 PB of data that needs to be shared with users in 1,000 AWS accounts. The data is proprietary, so some of it should only be available to users with specific job types. Some of the data is used for throughput-intensive workloads such as simulations. The number of accounts changes frequently because of new initiatives, acquisitions, and divestitures.

A solutions architect has been asked to design a system that will allow for sharing data for use in AWS with all of the employees in the company.

Which approach will allow for secure data sharing in a scalable way?

- A. Store the data in a single Amazon S3 bucket. Create an IAM role for every combination of job type and business unit that allows for appropriate read/write access based on object prefixes in the S3 bucket. The roles should have trust policies that allow the business unit's AWS accounts to assume their roles. Use IAM in each business unit's AWS account to prevent them from assuming roles for a different job type. Users get credentials to access the data by using AssumeRole from their business unit's AWS account. Users can then use those credentials with an S3 client.
- B. Store the data in a single Amazon S3 bucket. Write a bucket policy that uses conditions to grant read and write access where appropriate based on each user's business unit and job type.
- C. Determine the business unit with the AWS account accessing the bucket and the job type with a prefix in the IAM user's name. Users can access data by using IAM credentials from their business unit's AWS account with an S3 client.
- D. Store the data in a series of Amazon S3 buckets. Create an application running on Amazon EC2 that is integrated with the company's identity provider (IdP) that authenticates users and allows them to download or upload data through the application. The application uses the business unit and job type information in the IdP to control what users can upload and download through the application. The users can access the data through the application's API.
- E. Store the data in a series of Amazon S3 buckets. Create an AWS STS token vending machine that is integrated with the company's identity provider (IdP). When a user logs in, have the token vending machine attach an IAM policy that assumes the role that limits the user's access and/or upload only the data the user is authorized to access. Users can get credentials by authenticating to the token vending machine's website or API and then use those credentials with an S3 client.
- F. D.

Answer: E

NEW QUESTION 190

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual SAP-C02 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the SAP-C02 Product From:

<https://www.2passeasy.com/dumps/SAP-C02/>

Money Back Guarantee

SAP-C02 Practice Exam Features:

- * SAP-C02 Questions and Answers Updated Frequently
- * SAP-C02 Practice Questions Verified by Expert Senior Certified Staff
- * SAP-C02 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * SAP-C02 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year