

Exam Questions SAP-C02

AWS Certified Solutions Architect - Professional

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



NEW QUESTION 1

- (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 2

- (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 VPC
- 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 VPC
- 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 VPC
- H. Create a Route 53 outbound resolver in the shared services VPC
- 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 VPC
- K. Create a Route 53 inbound resolver in the shared services VPC
- 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 3

- (Exam Topic 1)

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

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

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

P. Configure a bucket policy on the logging bucket that disallows uploads of non-encrypted data and requires that the encryption source be AWS KMS.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/security/how-to-byok-bring-your-own-key-to-aws-kms-for-less-than-15-00-a-yea>
<https://docs.aws.amazon.com/kms/latest/developerguide/importing-keys-create-cmk.html>

NEW QUESTION 4

- (Exam Topic 1)

An enterprise runs 103 line-of-business applications on virtual machines in an on-premises data center. Many of the applications are simple PHP, Java, or Ruby web applications, are no longer actively developed, and serve little traffic.

Which approach should be used to migrate these applications to AWS with the LOWEST infrastructure costs?

- A. Deploy the applications to single-instance AWS Elastic Beanstalk environments without a load balancer.
- B. Use AWS SMS to create AMIs for each virtual machine and run them in Amazon EC2.
- C. Convert each application to a Docker image and deploy to a small Amazon ECS cluster behind an Application Load Balancer.
- D. Use VM Import/Export to create AMIs for each virtual machine and run them in single-instance AWS Elastic Beanstalk environments by configuring a custom image.

Answer: C

NEW QUESTION 5

- (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-A
- 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 over head and automatic scaling

NEW QUESTION 6

- (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 7

- (Exam Topic 1)

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

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

Which set of steps should the solutions architect take?

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

Answer: C

Explanation:

Combination of On-Demand and Spot Instances + DynamoDB.

NEW QUESTION 8

- (Exam Topic 1)

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

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

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

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

Answer: D

NEW QUESTION 9

- (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 ALB
- C. Set one of the origins as primary.
- D. Create an Amazon Route 53 health check for each ALB
- E. Create a Route 53 failover routing record pointing to the two ALBs
- F. Set the Evaluate Target Health value to Yes.
- G. Create two Amazon CloudFront distributions, each with one ALB as the origin
- H. Create an Amazon Route 53 failover routing record pointing to the two CloudFront distributions
- I. Set the Evaluate Target Health value to Yes.
- J. Create an Amazon Route 53 health check for each ALB
- K. Create a Route 53 latency alias record pointing to the two ALBs
- 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 10

- (Exam Topic 1)

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

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

Answer: B

NEW QUESTION 10

- (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 14

- (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 destinatio
- B. Choose to send logs to an Amazon S3 bucket.
- C. Enable AWS CloudTrail loggin
- D. Specify an Amazon S3 bucket as the destination for the logs.
- E. Use Amazon Athena to query the fogs in the Amazon S3 bucket for recipient, subject, and time sent.
- F. Create an Amazon CloudWatch log grou
- 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 17

- (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 19

- (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 24

- (Exam Topic 1)

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

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

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

Answer: ACF

Explanation:

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

NEW QUESTION 27

- (Exam Topic 1)

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

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

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

- A. Store the data in an Amazon Aurora Serverless database.
- B. Serve the data through a Network Load Balancer (NLB). Authenticate users using the NLB with credentials stored in AWS Secrets Manager.

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

Answer: C

Explanation:

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

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

NEW QUESTION 28

- (Exam Topic 1)

A company has a new application that needs to run on five Amazon EC2 instances in a single AWS Region. The application requires high-throughput, low-latency network connections between all of the EC2 instances where the application will run. There is no requirement for the application to be fault tolerant. Which solution will meet these requirements?

- A. Launch five new EC2 instances into a cluster placement group
- B. Ensure that the EC2 instance type supports enhanced networking.
- C. Launch five new EC2 instances into an Auto Scaling group in the same Availability Zone
- D. Attach an extra elastic network interface to each EC2 instance.
- E. Launch five new EC2 instances into a partition placement group
- F. Ensure that the EC2 instance type supports enhanced networking.
- G. Launch five new EC2 instances into a spread placement group
- H. Attach an extra elastic network interface to each EC2 instance.

Answer: A

Explanation:

When you launch EC2 instances in a cluster they benefit from performance and low latency. No redundancy though as per the question <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>.

NEW QUESTION 33

- (Exam Topic 1)

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

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

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

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

Answer: B

NEW QUESTION 38

- (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 43

- (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 that they 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 companies
- 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 stacks
- K. Incorporate the AWS CDK as a CodeBuild job in CodePipeline.

Answer: D

NEW QUESTION 48

- (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 50

- (Exam Topic 1)

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

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

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

Answer: C

Explanation:

"The company needs to reduce the cost and operational complexity for storing and serving this data. Which solution meets these requirements in the MOST cost-

effective manner?" EMR storage is ephemeral. The company has 100TB that need to persist, they would have to use EMRFS to backup to S3 anyway.
<https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-plan-storage.html>

100TB

EBS - 8.109\$ S3 - 2.355\$

You have saved 5.752\$

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

NEW QUESTION 51

- (Exam Topic 1)

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

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

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

Answer: BCE

Explanation:

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

NEW QUESTION 55

- (Exam Topic 1)

A company has many services running in its on-premises data center. The data center is connected to AWS using AWS Direct Connect (DX) and an IPsec VPN. The service data is sensitive and connectivity cannot traverse the internet. The company wants to expand into a new market segment and begin offering its services to other companies that are using AWS.

Which solution will meet these requirements?

- A. Create a VPC Endpoint Service that accepts TCP traffic, host it behind a Network Load Balancer, and make the service available over DX.
- B. Create a VPC Endpoint Service that accepts HTTP or HTTPS traffic, host it behind an Application Load Balancer, and make the service available over DX.
- C. Attach an internet gateway to the VPC
- D. and ensure that network access control and security group rules allow the relevant inbound and outbound traffic.
- E. Attach a NAT gateway to the VPC
- F. and ensure that network access control and security group rules allow the relevant inbound and outbound traffic.

Answer: A

NEW QUESTION 59

- (Exam Topic 1)

An online retail company hosts its stateful web-based application and MySQL database in an on-premises data center on a single server. The company wants to increase its customer base by conducting more marketing campaigns and promotions. In preparation, the company wants to migrate its application and database to AWS to increase the reliability of its architecture.

Which solution should provide the HIGHEST level of reliability?

- A. Migrate the database to an Amazon RDS MySQL Multi-AZ DB instance
- B. Deploy the application in an Auto Scaling group on Amazon EC2 instances behind an Application Load Balance
- C. Store sessions in Amazon Neptune.
- D. Migrate the database to Amazon Aurora MySQL
- E. Deploy the application in an Auto Scaling group on Amazon EC2 instances behind an Application Load Balance
- F. Store sessions in an Amazon ElastiCache for Redis replication group.
- G. Migrate the database to Amazon DocumentDB (with MongoDB compatibility). Deploy the application in an Auto Scaling group on Amazon EC2 instances behind a Network Load Balance
- H. Store sessions in Amazon Kinesis Data Firehose.
- I. Migrate the database to an Amazon RDS MariaDB Multi-AZ DB instance
- J. Deploy the application in an Auto Scaling group on Amazon EC2 instances behind an Application Load Balance
- K. Store sessions in Amazon ElastiCache for Memcached.

Answer: B

NEW QUESTION 61

- (Exam Topic 1)

To abide by industry regulations, a solutions architect must design a solution that will store a company's critical data in multiple public AWS Regions, including in the United States, where the company's headquarters is located. The solutions architect is required to provide access to the data stored in AWS to the company's global WAN network. The security team mandates that no traffic accessing this data should traverse the public internet.

How should the solutions architect design a highly available solution that meets the requirements and is cost-effective?

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

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

Answer: D

Explanation:

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

NEW QUESTION 63

- (Exam Topic 1)

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

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

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 68

- (Exam Topic 1)

A company has developed a single-page web application in JavaScript. The source code is stored in a single Amazon S3 bucket in the us-east-1 Region. The company serves the web application to a global user base through Amazon CloudFront.

The company wants to experiment with two versions of the website without informing application users. Each version of the website will reside in its own S3 bucket. The company wants to determine which version is most successful in marketing a new product.

The solution must send application users that are based in Europe to the new website design. The solution must send application users that are based in the United States to the current website design. However, some exceptions exist. The company needs to be able to redirect specific users to the new website design, regardless of the users' location.

Which solution meets these requirements?

- A. Configure two CloudFront distribution
- B. Configure a geolocation routing policy in Amazon Route 53 to route traffic to the appropriate CloudFront endpoint based on the location of clients.
- C. Configure a single CloudFront distributio
- D. Create a behavior with different paths for each version of the sit
- E. Configure Lambda@Edge on the default path to generate redirects and send the client to the correct version of the website.
- F. Configure a single CloudFront distributio
- G. Configure an alternate domain name on the distribution. Configure two behaviors to route users to the different S3 origins based on the domain name that the client uses in the HTTP request.
- H. Configure a single CloudFront distribution with Lambda@Edg
- I. Use Lambda@Edge to send user requests to different origins based on request attributes.

Answer: A

NEW QUESTION 69

- (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 Site-to-Site VPN connection from the business unit VPCs to the shared VPC
 N. Configure VPC routing tables to send traffic to the VPN connection.

Answer: B

Explanation:

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

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

NEW QUESTION 71

- (Exam Topic 1)

A company is launching a new web application on Amazon EC2 instances. Development and production workloads exist in separate AWS accounts. According to the company's security requirements, only automated configuration tools are allowed to access the production account. The company's security team wants to receive immediate notification if any manual access to the production AWS account or EC2 instances occurs. Which combination of actions should a solutions architect take in the production account to meet these requirements? (Select THREE.)

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

Answer: CDE

NEW QUESTION 73

- (Exam Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 75

- (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 distribution.
- 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 distribution.
- 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 76

- (Exam Topic 1)

A company is creating a REST API to share information with six of its partners based in the United States. The company has created an Amazon API Gateway Regional endpoint. Each of the six partners will access the API once per day to post daily sales figures. After initial deployment, the company observes 1.000 requests per second originating from 500 different IP addresses around the world. The company believes this traffic is originating from a botnet and wants to secure its API while minimizing cost. Which approach should the company take to secure its API?

- A. Create an Amazon CloudFront distribution with the API as the origin
- B. Create an AWS WAF web ACL with a rule to block clients that submit more than five requests per day
- C. Associate the web ACL with the CloudFront distribution
- D. Configure CloudFront with an origin access identity (OAI) and associate it with the distribution
- E. Configure API Gateway to ensure only the OAI can execute the POST method.
- F. Create an Amazon CloudFront distribution with the API as the origin
- G. Create an AWS WAF web ACL with a rule to block clients that submit more than five requests per day
- H. Associate the web ACL with the CloudFront distribution
- I. Add a custom header to the CloudFront distribution populated with an API key
- J. Configure the API to require an API key on the POST method.
- K. Create an AWS WAF web ACL with a rule to allow access to the IP addresses used by the six partners. Associate the web ACL with the API
- L. Create a resource policy with a request limit and associate it with the API
- M. Configure the API to require an API key on the POST method.
- N. Associate the web ACL with the API
- O. Create a usage plan with a request limit and associate it with the API
- P. Create an API key and add it to the usage plan.

Answer: D

Explanation:

"A usage plan specifies who can access one or more deployed API stages and methods—and also how much and how fast they can access them. The plan uses API keys to identify API clients and meters access to the associated API stages for each key. It also lets you configure throttling limits and quota limits that are enforced on individual client API keys."

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-usage-plans.html>

NEW QUESTION 81

- (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 86

- (Exam Topic 1)

A company is serving files to its customers through an SFTP server that is accessible over the internet. The SFTP server is running on a single Amazon EC2 instance with an Elastic IP address attached. Customers connect to the SFTP server through its Elastic IP address and use SSH for authentication. The EC2 instance also has an attached security group that allows access from all customer IP addresses.

A solutions architect must implement a solution to improve availability, minimize the complexity of infrastructure management, and minimize the disruption to customers who access files. The solution must not change the way customers connect.

Which solution will meet these requirements?

- A. Disassociate the Elastic IP address from the EC2 instance
- B. Create an Amazon S3 bucket to be used for SFTP file hosting
- C. Create an AWS Transfer Family server. Configure the Transfer Family server with a publicly accessible endpoint. Associate the SFTP Elastic IP address with the new endpoint. Point the Transfer Family server to the S3 bucket.
- D. Sync all files from the SFTP server to the S3 bucket.
- E. Disassociate the Elastic IP address from the EC2 instance
- F. Create an Amazon S3 bucket to be used for SFTP file hosting
- G. Create an AWS Transfer Family server
- H. Configure the Transfer Family server with a VPC-hosted internet-facing endpoint
- I. Associate the SFTP Elastic IP address with the new endpoint
- K. Attach the security group with customer IP addresses to the new endpoint
- L. Point the Transfer Family server to the S3 bucket. Sync all files from the SFTP server to the S3 bucket.
- M. Disassociate the Elastic IP address from the EC2 instance
- N. Create a new Amazon Elastic File System (Amazon EFS) file system to be used for SFTP file hosting
- O. Create an AWS Fargate task definition to run an SFTP server
- P. Specify the EFS file system as a mount in the task definition
- Q. Create a Fargate service by using the task definition, and place a Network Load Balancer (NLB) in front of the service. When configuring the service, attach the security group with customer IP addresses to the tasks that run the SFTP server
- R. Associate the Elastic IP address with the NLB
- S. Sync all files from the SFTP server to the S3 bucket.

- T. Disassociate the Elastic IP address from the EC2 instance
- . Create a multi-attach Amazon Elastic Block Store (Amazon EBS) volume to be used for SFTP file hosting
- . Create a Network Load Balancer (NLB) with the Elastic IP address attached
- . Create an Auto Scaling group with EC2 instances that run an SFTP server. Define in the Auto Scaling group that instances that are launched should attach the new multi-attach EBS volume. Configure the Auto Scaling group to automatically add instances behind the NLB. Configure the Auto Scaling group to use the security group that allows customer IP addresses for the EC2 instances that the Auto Scaling group launches
- . Sync all files from the SFTP server to the new multi-attach EBS volume.

Answer: B

Explanation:

<https://docs.aws.amazon.com/transfer/latest/userguide/create-server-in-vpc.html> <https://aws.amazon.com/premiumsupport/knowledge-center/aws-sftp-endpoint-type/>

NEW QUESTION 89

- (Exam Topic 1)

A solution architect needs to deploy an application on a fleet of Amazon EC2 instances. The EC2 instances run in private subnets in an Auto Scaling group. The application is expected to generate logs at a rate of 100 MB each second on each of the EC2 instances.

The logs must be stored in an Amazon S3 bucket so that an Amazon EMR cluster can consume them for further processing. The logs must be quickly accessible for the first 90 days and should be retrievable within 48 hours thereafter.

What is the MOST cost-effective solution that meets these requirements?

- A. Set up an S3 copy job to write logs from each EC2 instance to the S3 bucket with S3 Standard storage. Use a NAT instance within the private subnets to connect to Amazon S3. Create S3 Lifecycle policies to move logs that are older than 90 days to S3 Glacier.
- B. Set up an S3 sync job to copy logs from each EC2 instance to the S3 bucket with S3 Standard storage. Use a gateway VPC endpoint for Amazon S3 to connect to Amazon S3. Create S3 Lifecycle policies to move logs that are older than 90 days to S3 Glacier Deep Archive.
- C. Set up an S3 batch operation to copy logs from each EC2 instance to the S3 bucket with S3 Standard storage. Use a NAT gateway with the private subnets to connect to Amazon S3. Create S3 Lifecycle policies to move logs that are older than 90 days to S3 Glacier Deep Archive.
- D. Set up an S3 sync job to copy logs from each EC2 instance to the S3 bucket with S3 Standard storage. Use a gateway VPC endpoint for Amazon S3 to connect to Amazon S3. Create S3 Lifecycle policies to move logs that are older than 90 days to S3 Glacier.

Answer: C

NEW QUESTION 93

- (Exam Topic 1)

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

- A. Create multiple read replicas and put them into an Auto Scaling group.
- B. Create multiple read replicas in different Availability Zones.
- C. Create an Amazon Route 53 hosted zone and a record set for each read replica with a TTL and a weighted routing policy.
- D. Create an Application Load Balancer (ALB) and put the read replicas behind the ALB.
- E. Configure an Amazon CloudWatch alarm to detect a failed read replica.
- 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 95

- (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 store.
- 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 Virtualization Format (OVF) format. Create an Amazon S3 bucket to store the image in the destination AWS Region.
- 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 file service to export a Common Internet File System (CIFS) share.
- F. Create a backup copy to the shared folder.
- G. Sign in to the AWS Management Console and create an AMI from the backup copy. Launch an EC2 instance that is based on the AMI.
- H. Create a managed-instance activation for a hybrid environment in AWS Systems Manager.
- I. Download and install Systems Manager Agent on the on-premises VM. Register the VM with Systems Manager to be a managed instance. Use AWS Backup to create a snapshot of the VM and create an AMI.
- J. Launch an EC2 instance that is based on the AMI.

Answer: B

Explanation:

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

- Export an OVF Template
- Create / use an Amazon S3 bucket for storing the exported images. The bucket must be in the Region where you want to import your VMs.

- Create an IAM role named vmimport.
- You'll use AWS CLI to run the import commands. <https://aws.amazon.com/premiumsupport/knowledge-center/import-instances/>

NEW QUESTION 99

- (Exam Topic 1)

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

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

Answer: BCF

Explanation:

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

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

NEW QUESTION 101

- (Exam Topic 1)

An AWS customer has a web application that runs on premises. The web application fetches data from a third-party API that is behind a firewall. The third party accepts only one public CIDR block in each client's allow list. The customer wants to migrate their web application to the AWS Cloud. The application will be hosted on a set of Amazon EC2 instances behind an Application Load Balancer (ALB) in a VPC. The ALB is located in public subnets. The EC2 instances are located in private subnets. NAT gateways provide internet access to the private subnets. How should a solutions architect ensure that the web application can continue to call the third-party API after the migration?

- A. Associate a block of customer-owned public IP addresses to the VPC
- B. Enable public IP addressing for public subnets in the VPC.
- C. Register a block of customer-owned public IP addresses in the AWS account
- D. Create Elastic IP addresses from the address block and assign them to the NAT gateways in the VPC.
- E. Create Elastic IP addresses from the block of customer-owned IP addresses
- F. Assign the static Elastic IP addresses to the ALB.
- G. Register a block of customer-owned public IP addresses in the AWS account
- H. Set up AWS Global Accelerator to use Elastic IP addresses from the address block
- I. Set the ALB as the accelerator endpoint.

Answer: B

Explanation:

When EC2 instances reach third-party API through internet, their private IP addresses will be masked by NAT Gateway public IP address. <https://aws.amazon.com/blogs/networking-and-content-delivery/introducing-bring-your-own-ip-byoip-for-amaz>

NEW QUESTION 106

- (Exam Topic 1)

A North American company with headquarters on the East Coast is deploying a new web application running on Amazon EC2 in the us-east-1 Region. The application should dynamically scale to meet user demand and maintain resiliency. Additionally, the application must have disaster recovery capabilities in an active-passive configuration with the us-west-1 Region.

Which steps should a solutions architect take after creating a VPC in the us-east-1 Region?

- A. Create a VPC in the us-west-1 Region
- B. Use inter-Region VPC peering to connect both VPCs
- C. Deploy an Application Load Balancer (ALB) spanning multiple Availability Zones (AZs) to the VPC in the us-east-1 Region
- D. Deploy EC2 instances across multiple AZs in each Region as part of an Auto Scaling group spanning both VPCs and served by the ALB.
- E. Deploy an Application Load Balancer (ALB) spanning multiple Availability Zones (AZs) to the VPC in the us-east-1 Region
- F. Deploy EC2 instances across multiple AZs as part of an Auto Scaling group served by the ALB
- G. Deploy the same solution to the us-west-1 Region. Create an Amazon Route 53 record set with a failover routing policy and health checks enabled to provide high availability across both Regions.
- H. Create a VPC in the us-west-1 Region
- I. Use inter-Region VPC peering to connect both VPCs. Deploy an Application Load Balancer (ALB) that spans both VPCs. Deploy EC2 instances across multiple Availability Zones as part of an Auto Scaling group in each VPC served by the ALB.
- J. Create an Amazon Route 53 record that points to the ALB.
- K. Deploy an Application Load Balancer (ALB) spanning multiple Availability Zones (AZs) to the VPC in the us-east-1 Region
- L. Deploy EC2 instances across multiple AZs as part of an Auto Scaling group served by the ALB
- M. Deploy the same solution to the us-west-1 Region
- N. Create separate Amazon Route 53 records in each Region that point to the ALB in the Region
- O. Use Route 53 health checks to provide high availability across both Regions.

Answer: B

Explanation:

A new web application in an active-passive DR mode. a Route 53 record set with a failover routing policy.

NEW QUESTION 111

- (Exam Topic 1)

A company requires that all internal application connectivity use private IP addresses. To facilitate this policy, a solutions architect has created interface endpoints to connect to AWS public services. Upon testing, the solutions architect notices that the service names are resolving to public IP addresses, and that internal services cannot connect to the interface endpoints.

Which step should the solutions architect take to resolve this issue?

- A. Update the subnet route table with a route to the interface endpoint.
- B. Enable the private DNS option on the VPC attributes.
- C. Configure the security group on the interface endpoint to allow connectivity to the AWS services.
- D. Configure an Amazon Route 53 private hosted zone with a conditional forwarder for the internal application.

Answer: C

Explanation:

<https://docs.aws.amazon.com/vpc/latest/privatelink/vpce-interface.html>

NEW QUESTION 113

- (Exam Topic 1)

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

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

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

Answer: ACF

Explanation:

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

NEW QUESTION 114

- (Exam Topic 1)

A company runs a popular public-facing ecommerce website. Its user base is growing quickly from a local market to a national market. The website is hosted in an on-premises data center with web servers and a MySQL database. The company wants to migrate its workload to AWS. A solutions architect needs to create a solution to:

- Improve security
- Improve reliability
- Improve availability
- Reduce latency
- Reduce maintenance

Which combination of steps should the solutions architect take to meet these requirements? (Select THREE.)

- A. Use Amazon EC2 instances in two Availability Zones for the web servers in an Auto Scaling group behind an Application Load Balancer.
- B. Migrate the database to a Multi-AZ Amazon Aurora MySQL DB cluster.
- C. Use Amazon EC2 instances in two Availability Zones to host a highly available MySQL database cluster.
- D. Host static website content in Amazon S3. Use S3 Transfer Acceleration to reduce latency while serving webpage
- E. Use AWS WAF to improve website security.
- F. Host static website content in Amazon S3. Use Amazon CloudFront to reduce latency while serving webpage
- G. Use AWS WAF to improve website security
- H. Migrate the database to a single-AZ Amazon RDS for MySQL DB instance.

Answer: ABE

NEW QUESTION 118

- (Exam Topic 1)

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

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

Which solution will meet these requirements?

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

- O. Create an Application Load Balancer (ALB), and attach it to the endpoint group
- P. Point the ALB to the on-premises server
- Q. Offload the DNS querying to AWS to handle AWS Global Accelerator traffic.

Answer: D

NEW QUESTION 119

- (Exam Topic 1)

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 124

- (Exam Topic 1)

A company is planning on hosting its e-commerce platform on AWS using a multi-tier web application designed for a NoSQL database. The company plans to use the us-west-2 Region as its primary Region. The company wants to ensure that copies of the application and data are available in a second Region, us-west-1, for disaster recovery. The company wants to keep the time to fail over as low as possible. Failing back to the primary Region should be possible without administrative interaction after the primary service is restored.

Which design should the solutions architect use?

- A. Use AWS CloudFormation StackSets to create the stacks in both Regions with Auto Scaling groups for the web and application tier
- B. Asynchronously replicate static content between Regions using Amazon S3 cross-Region replication
- C. Use an Amazon Route 53 DNS failover routing policy to direct users to the secondary site in us-west-1 in the event of an outage
- D. Use Amazon DynamoDB global tables for the database tier.
- E. Use AWS CloudFormation StackSets to create the stacks in both Regions with Auto Scaling groups for the web and application tier
- F. Asynchronously replicate static content between Regions using Amazon S3 cross-Region replication
- G. Use an Amazon Route 53 DNS failover routing policy to direct users to the secondary site in us-west-1 in the event of an outage
- H. Deploy an Amazon Aurora global database for the database tier.
- I. Use AWS Service Catalog to deploy the web and application servers in both Regions
- J. Asynchronously replicate static content between the two Regions using Amazon S3 cross-Region replication
- K. Use Amazon Route 53 health checks to identify a primary Region failure and update the public DNS entry listing to the secondary Region in the event of an outage
- L. Use Amazon RDS for MySQL with cross-Region replication for the database tier.
- M. Use AWS CloudFormation StackSets to create the stacks in both Regions using Auto Scaling groups for the web and application tier
- N. Asynchronously replicate static content between Regions using Amazon S3 cross-Region replication
- O. Use Amazon CloudFront with static files in Amazon S3, and multi-Region origins for the front-end web tier
- P. Use Amazon DynamoDB tables in each Region with scheduled backups to Amazon S3.

Answer: A

NEW QUESTION 128

- (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 NAT Gateway-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 129

- (Exam Topic 1)

A developer reports receiving an Error 403: Access Denied message when they try to download an object from an Amazon S3 bucket. The S3 bucket is accessed using an S3 endpoint inside a VPC, and is encrypted with an AWS KMS key. A solutions architect has verified that (he developer is assuming the correct IAM role in the account that allows the object to be downloaded. The S3 bucket policy and the NACL are also valid.

Which additional step should the solutions architect take to troubleshoot this issue?

- A. Ensure that blocking all public access has not been enabled in the S3 bucket.
- B. Verify that the IAM role has permission to decrypt the referenced KMS key.
- C. Verify that the IAM role has the correct trust relationship configured.
- D. Check that local firewall rules are not preventing access to the S3 endpoint.

Answer: B

NEW QUESTION 131

- (Exam Topic 1)

A company has a policy that all Amazon EC2 instances that are running a database must exist within the same subnets in a shared VPC. Administrators must follow security compliance requirements and are not allowed to directly log in to the shared account. All company accounts are members of the same organization in AWS Organizations. The number of accounts will rapidly increase as the company grows.

A solutions architect uses AWS Resource Access Manager to create a resource share in the shared account. What is the MOST operationally efficient configuration to meet these requirements?

- A. Add the VPC to the resource share
- B. Add the account IDs as principals
- C. Add all subnets within the VPC to the resource share
- D. Add the account IDs as principals
- E. Add all subnets within the VPC to the resource share
- F. Add the organization as a principal.
- G. Add the VPC to the resource share
- H. Add the organization as a principal

Answer: C

Explanation:

<https://docs.aws.amazon.com/ram/latest/userguide/getting-started-sharing.html#getting-started-sharing-create> To restrict resource sharing to only principals in your organization, choose Allow sharing with principals in your organization only.

<https://docs.aws.amazon.com/ram/latest/userguide/ram-ug.pdf>

NEW QUESTION 136

- (Exam Topic 1)

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

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

Which solution will satisfy these requirements with minimal operational overhead?

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

Answer: A

Explanation:

SMS can handle migrating the data volumes:

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

NEW QUESTION 141

- (Exam Topic 1)

A company is running a web application with On-Demand Amazon EC2 instances in Auto Scaling groups that scale dynamically based on custom metrics. After extensive testing, the company determines that the m5.2xlarge instance size is optimal for the workload. Application data is stored in db.r4.4xlarge Amazon RDS instances that are confirmed to be optimal. The traffic to the web application spikes randomly during the day.

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

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

Answer: B

Explanation:

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

NEW QUESTION 143

- (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 Mbps internet connection.

Which migration strategy meets these requirements?

- A. Create an AWS Database Migration Service (AWS DMS) replication instance.
- 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 bucket.
- E. Copy the data into the Amazon Redshift cluster.
- 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 146

- (Exam Topic 1)

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

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

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 147

- (Exam Topic 1)

A company is running an application on Amazon EC2 instances in three environments; development, testing, and production. The company uses AMIs to deploy the EC2 instances. The company builds the AMIs by using custom deployment scripts and infrastructure orchestration tools for each release in each environment. The company is receiving errors in its deployment process. Errors appear during operating system package downloads and during application code installation from a third-party Git hosting service. The company needs deployments to become more reliable across all environments. Which combination of steps will meet these requirements? (Select THREE).

- A. Mirror the application code to an AWS CodeCommit Git repositior
- B. Use the repository to build EC2 AMIs.
- C. Produce multiple EC2 AMI
- D. one for each environment, for each release.
- E. Produce one EC2 AMI for each release for use across all environments.
- F. Mirror the application code to a third-party Git repository that uses Amazon S3 storag
- G. Use therepository for deployment.
- H. Replace the custom scripts and tools with AWS CodeBuil
- I. Update the infrastructure deployment process to use EC2 Image Builder.

Answer: ACE

NEW QUESTION 151

- (Exam Topic 1)

A company is hosting a single-page web application in the AWS Cloud. The company is using Amazon CloudFront to reach its goal audience. The CloudFront distribution has an Amazon S3 bucket that is configured as its origin. The static files for the web application are stored in this S3 bucket. The company has used a simple routing policy to configure an Amazon Route 53 A record The record points to the CloudFront distribution The company wants to use a canary deployment release strategy for new versions of the application. What should a solutions architect recommend to meet these requirements?

- A. Create a second CloudFront distribution for the new version of the applicatio
- B. Update the Route 53 record to use a weighted routing policy.
- C. Create a Lambda@Edge functio
- D. Configure the function to implement a weighting algorithm and rewrite the URL to direct users to a new version of the application.
- E. Create a second S3 bucket and a second CloudFront origin for the new S3 bucket Create a CloudFrontorigin group that contains both origins Configure origin weighting for the origin group.
- F. Create two Lambda@Edge function
- G. Use each function to serve one of the application versions Set up a CloudFront weighted Lambda@Edge invocation policy

Answer: A

NEW QUESTION 153

- (Exam Topic 2)

A company is in the process of implementing AWS Organizations to constrain its developers to use only Amazon EC2, Amazon S3 and Amazon DynamoDB. The developers account resides In a dedicated organizational unit (OU). The solutions architect has implemented the following SCP on the developers account:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowEC2",
      "Effect": "Allow",
      "Action": "ec2:*",
      "Resource": "*"
    },
    {
      "Sid": "AllowDynamoDB",
      "Effect": "Allow",
      "Action": "dynamodb:*",
      "Resource": "*"
    },
    {
      "Sid": "AllowS3",
      "Effect": "Allow",
      "Action": "s3:*",
      "Resource": "*"
    }
  ]
}
```

When this policy is deployed, IAM users in the developers account are still able to use AWS services that are not listed in the policy. What should the solutions architect do to eliminate the developers' ability to use services outside the scope of this policy?

- A. Create an explicit deny statement for each AWS service that should be constrained
- B. Remove the Full AWS Access SCP from the developer account's OU
- C. Modify the Full AWS Access SCP to explicitly deny all services
- D. Add an explicit deny statement using a wildcard to the end of the SCP

Answer: B

NEW QUESTION 158

- (Exam Topic 2)

A company is migrating its marketing website and content management system from an on-premises data center to AWS. The company wants the AWS application to be deployed in a VPC with Amazon EC2 instances used for the web servers and an Amazon RDS instance for the database.

The company has a runbook document that describes the installation process of the on-premises system. The company would like to base the AWS system on the processes referenced in the runbook document. The runbook document describes the installation and configuration of the operating systems, network settings, the website, and content management system software on the servers. After the migration is complete, the company wants to be able to make changes quickly to take advantage of other AWS features.

How can the application and environment be deployed and automated in AWS, while allowing for future changes?

- A. Update the runbook to describe how to create the VPC
- B. the EC2 instances and the RDS instance for the application by using the AWS Console. Make sure that the rest of the steps in the runbook are updated to reflect any changes that may come from the AWS migration
- C. Write a Python script that uses the AWS API to create the VPC
- D. the EC2 instances and the RDS instance for the application. Write shell scripts that implement the rest of the steps in the runbook. Have the Python script copy and run the shell scripts on the newly created instances to complete the installation
- E. Write an AWS CloudFormation template that creates the VPC, the EC2 instances, and the RDS instance for the application. Ensure that the rest of the steps in the runbook are updated to reflect any changes that may come from the AWS migration
- F. Write an AWS CloudFormation template that creates the VPC, the EC2 instances, and the RDS instance for the application. Include EC2 user data in the AWS CloudFormation template to install and configure the software.

Answer: D

NEW QUESTION 159

- (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 164

- (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 VPC
- G. configure an end-to-end VPC, and advertise the routes for 172.50.0.0/16

Answer: C

NEW QUESTION 165

- (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 train 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 168

- (Exam Topic 2)

A company that runs applications on AWS recently subscribed to a new software-as-a-service (SaaS) data vendor. The vendor provides the data by way of a

REST API that the vendor hosts in its AWS environment. The vendor offers multiple options for connectivity to the API and is working with the company to find the best way to connect. The company's AWS account does not allow outbound internet access from its AWS environment. The vendor's services run on AWS in the same AWS Region as the company's applications. A solutions architect must implement connectivity to the vendor's API so that the API is highly available in the company's VPC. Which solution will meet these requirements?

- A. Connect to the vendor's public API address for the data service.
- B. Connect to the vendor by way of a VPC peering connection between the vendor's VPC and the company's VPC.
- C. Connect to the vendor by way of a VPC endpoint service that uses AWS PrivateLink.
- D. Connect to a public bastion host that the vendor provides. Tunnel the API traffic.

Answer: C

NEW QUESTION 171

- (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 172

- (Exam Topic 2)

A company is using an Amazon CloudFront distribution to distribute both static and dynamic content from a web application running behind an Application Load Balancer. The web application requires user authorization and session tracking for dynamic content. The CloudFront distribution has a single cache behavior configured to forward the Authorization, Host, and Agent HTTP allow list headers and a session cookie to the origin. All other cache behavior settings are set to their default value.

A valid ACM certificate is applied to the CloudFront distribution with a matching CNAME in the distribution settings. The ACM certificate is also applied to the HTTPS listener for the Application Load Balancer. The CloudFront origin protocol policy is set to HTTPS only. Analysis of the cache statistics report shows that the miss rate for this distribution is very high.

What can the solutions architect do to improve the cache hit rate for this distribution without causing the SSL/TLS handshake between CloudFront and the Application Load Balancer to fail?

- A. Create two cache behaviors for static and dynamic content. Remove the user-Agent and Host HTTP headers from the allow list headers section on both of the cache behaviors. Remove the session cookie from the allow list cookies section and the Authorization HTTP header from the allow list headers section for cache behavior configured for static content.
- B. Remove the user-Agent and Authorization HTTP headers from the allow list headers section of the cache behavior.
- C. Then update the cache behavior to use resigned cookies for authorization.
- D. Remove the Host HTTP header from the allow list headers section and remove the session cookie from the allow list cookies section for the default cache behavior. Enable automatic object compression and use Lambda@Edge viewer request events for user authorization.
- E. Create two cache behaviors for static and dynamic content. Remove the User-Agent HTTP header from the allow list headers section on both of the cache behaviors. Remove the session cookie from the allow list cookies section and the Authorization HTTP header from the allow list headers section for cache behavior configured for static content.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/understanding-the-cache-key.html> Removing the host header will result in failed flow between CloudFront and ALB, because they have same certificate.

NEW QUESTION 173

- (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 applicatio
- 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 databas
- 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 177

- (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 180

- (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 184

- (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 187

- (Exam Topic 2)

A company wants to use Amazon Workspaces in combination with thin client devices to replace aging desktops. Employees use the desktops to access applications that work with clinical trial data. Corporate security policy states that access to the applications must be restricted to only company branch office locations. The company is considering adding an additional branch office in the next 6 months. Which solution meets these requirements with the MOST operational efficiency?

- A. Create an IP access control group rule with the list of public addresses from the branch offices. Associate the IP access control group with the Workspaces directory.
- B. Use AWS Firewall Manager to create a web ACL rule with an IPSet with the list of public addresses from the branch office locations. Associate the web ACL with the Workspaces directory.
- C. Use AWS Certificate Manager (ACM) to issue trusted device certificates to the machines deployed in the branch office locations. Enable restricted access on the Workspaces directory.
- D. Create a custom Workspace image with Windows Firewall configured to restrict access to the public addresses of the branch offices. Use the image to deploy

the Workspaces.

Answer: C

NEW QUESTION 189

- (Exam Topic 2)

A company has an organization in AWS Organizations that has a large number of AWS accounts. One of the AWS accounts is designated as a transit account and has a transit gateway that is shared with all of the other AWS accounts. AWS Site-to-Site VPN connections are configured between all of the company's global offices and the transit account. The company has AWS Config enabled on all of its accounts.

The company's networking team needs to centrally manage a list of internal IP address ranges that belong to the global offices. Developers will reference this list to gain access to applications securely.

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

- A. Create a JSON file that is hosted in Amazon S3 and that lists all of the internal IP address ranges. Configure an Amazon Simple Notification Service (Amazon SNS) topic in each of the accounts that can be involved when the JSON file is updated.
- B. Subscribe an AWS Lambda function to the SNS topic to update all relevant security group rules with the updated IP address ranges.
- C. Create a new AWS Config managed rule that contains all of the internal IP address ranges. Use the rule to check the security groups in each of the accounts to ensure compliance with the list of IP address ranges.
- D. Configure the rule to automatically remediate any noncompliant security group that is detected.
- E. In the transit account, create a VPC prefix list with all of the internal IP address ranges.
- F. Use AWS Resource Access Manager to share the prefix list with all of the other accounts.
- G. Use the shared prefix list to configure security group rules in the other accounts.
- H. In the transit account, create a security group with all of the internal IP address ranges.
- I. Configure the security groups in the other accounts to reference the transit account's security group by using a nested security group reference of `*<transit-account-id>./sg-1a2b3c4d`.

Answer: C

NEW QUESTION 191

- (Exam Topic 2)

A retail company is running an application that stores invoice files in an Amazon S3 bucket and metadata about the files in an Amazon DynamoDB table. The application software runs in both us-east-1 and eu-west-1. The S3 bucket and DynamoDB table are in us-east-1. The company wants to protect itself from data corruption and loss of connectivity to either region.

Which option meets these requirements?

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

Answer: B

NEW QUESTION 196

- (Exam Topic 2)

A solutions architect uses AWS Organizations to manage several AWS accounts for a company. The full Organizations feature set is activated for the organization. All production AWS accounts exist under an OU that is named "production". Systems operators have full administrative privileges within these accounts by using IAM roles.

The company wants to ensure that security groups in all production accounts do not allow inbound traffic for TCP port 22. All noncompliant security groups must be remediated immediately, and no new rules that allow port 22 can be created.

Which solution will meet these requirements?

- A. Write an SCP that denies the `CreateSecurityGroup` action with a condition of `ec2:ingress rule with value 22`. Apply the SCP to the "production" OU.
- B. Configure an AWS CloudTrail trail for all accounts. Send CloudTrail logs to an Amazon S3 bucket in the Organizations management account.
- C. Configure an AWS Lambda function on the management account with permissions to assume a role in all production accounts to describe and modify security groups.
- D. Configure Amazon S3 to invoke the Lambda function on every `PutObject` event on the S3 bucket. Configure the Lambda function to analyze each CloudTrail event for noncompliant security group actions and to automatically remediate any issues.
- E. Create an Amazon EventBridge (Amazon CloudWatch Events) event bus in the Organizations management account.
- F. Create an AWS CloudFormation template to deploy configurations that send `CreateSecurityGroup` events to the event bus from all production accounts. Configure an AWS Lambda function in the management account with permissions to assume a role in all production accounts to describe and modify security groups.
- G. Configure the event bus to invoke the Lambda function. Configure the Lambda function to analyze each event for noncompliant security group actions and to automatically remediate any issues.
- H. Create an AWS CloudFormation template to turn on AWS Config. Activate the `INCOMING_SSH_DISABLED` AWS Config managed rule. Deploy an AWS Lambda function that will run based on AWS Config findings and will remediate noncompliant resources. Deploy the CloudFormation template by using a StackSet that is assigned to the "production" OU.
- I. Apply an SCP to the OU to deny modification of the resources that the CloudFormation template provisions.

Answer: D

NEW QUESTION 200

- (Exam Topic 2)

A solutions architect has been assigned to migrate a 50 TB Oracle data warehouse that contains sales data from on-premises to Amazon Redshift. Major updates to the sales data occur on the final calendar day of the month. For the remainder of the month, the data warehouse only receives minor daily updates and is primarily used for reading and reporting. Because of this, the migration process must start on the first day of the month and must be complete before the next set of

updates occur. This provides approximately 30 days to complete the migration and ensure that the minor daily changes have been synchronized with the Amazon Redshift data warehouse. Because the migration cannot impact normal business network operations, the bandwidth allocated to the migration for moving data over the internet is 50 Mbps. The company wants to keep data migration costs low.

Which steps will allow the solutions architect to perform the migration within the specified timeline?

- A. Install Oracle database software on an Amazon EC2 instance. Configure VPN connectivity between AWS and the company's data center. Configure the Oracle database running on Amazon EC2 to join the Oracle Real Application Clusters (RAC). When the Oracle database on Amazon EC2 finishes synchronizing, create an AWS DMS ongoing replication task to migrate the data from the Oracle database on Amazon EC2 to Amazon Redshift. Verify the data migration is complete and perform the cut over to Amazon Redshift.
- B. Create an AWS Snowball import job. Export a backup of the Oracle data warehouse. Copy the exported data to the Snowball device. Return the Snowball device to AWS. Create an Amazon RDS for Oracle database and restore the backup file to that RDS instance. Create an AWS DMS task to migrate the data from the RDS for Oracle database to Amazon Redshift. Copy daily incremental backups from Oracle in the data center to the RDS for Oracle database over the internet. Verify the data migration is complete and perform the cut over to Amazon Redshift.
- C. Install Oracle database software on an Amazon EC2 instance. To minimize the migration time, configure VPN connectivity between AWS and the company's data center by provisioning a 1 Gbps AWS Direct Connect connection. Configure the Oracle database running on Amazon EC2 to be a read replica of the data center Oracle database. Start the synchronization process between the company's on-premises data center and the Oracle database on Amazon EC2. When the Oracle database on Amazon EC2 is synchronized with the on-premises database, create an AWS DMS ongoing replication task from the Oracle database read replica that is running on Amazon EC2 to Amazon Redshift. Verify the data migration is complete and perform the cut over to Amazon Redshift.
- D. Create an AWS Snowball import job.
- E. Configure a server in the company's data center with an extraction agent.
- F. Use AWS SCT to manage the extraction agent and convert the Oracle schema to an Amazon Redshift schema.
- G. Create a new project in AWS SCT using the registered data extraction agent.
- H. Create a local task and an AWS DMS task in AWS SCT with replication of ongoing changes.
- I. Copy data to the Snowball device and return the Snowball device to AWS.
- J. Allow AWS DMS to copy data from Amazon S3 to Amazon Redshift.
- K. Verify that the data migration is complete and perform the cut over to Amazon Redshift.

Answer: D

Explanation:

Create an AWS Snowball import job. Configure a server in the company's data center with an extraction agent. Use AWS SCT to manage the extraction agent and convert the Oracle schema to an Amazon Redshift schema. Create a new project in AWS SCT using the registered data extraction agent. Create a local task and an AWS DMS task in AWS SCT with replication of ongoing changes. Copy data to the Snowball device and return the Snowball device to AWS. Allow AWS DMS to copy data from Amazon S3 to Amazon Redshift. Verify that the data migration is complete and perform the cut over to Amazon Redshift.
<https://aws.amazon.com/getting-started/hands-on/migrate-oracle-to-amazon-redshift/>

NEW QUESTION 202

.....

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