[2025-November-NewBraindump2go SAP-C02 Dumps PDF Free[Q175-Q206

2025/November Latest Braindump2go SAP-C02 Exam Dumps with PDF and VCE Free Updated Today! Following are some new Braindump2go SAP-C02 Real Exam Questions!QUESTION 175A company is developing a new service that will be accessed using TCP on a static port. A solutions architect must ensure that the service is highly available, has redundancy across Availability Zones, and is accessible using the DNS name my.service.com, which is publicly accessible. The service must use fixed address assignments so other companies can add the addresses to their allow lists.Assuming that resources are deployed in multiple Availability Zones in a single Region, which solution will meet these requirements?A. Create Amazon EC2 instances with an Elastic IP address for each instance. Create a Network Load Balancer (NLB) and expose the static TCP port. Register EC2 instances with the NLB. Create a new name server record set named my.service.com, and assign the Elastic IP addresses of the EC2 instances to the record set. Provide the Elastic IP addresses of the EC2 instances to the other companies to add to their allow lists.B. Create an Amazon ECS cluster and a service definition for the application. Create and assign public IP addresses for the ECS cluster. Create a Network Load Balancer (NLB) and expose the TCP port. Create a target group and assign the ECS cluster name to the NLB. Create a new A record set named my.service.com, and assign the public IP addresses of the ECS cluster to the record set. Provide the public IP addresses of the ECS cluster to the other companies to add to their allow lists.C. Create Amazon EC2 instances for the service. Create one Elastic IP address for each Availability Zone. Create a Network Load Balancer (NLB) and expose the assigned TCP port. Assign the Elastic IP addresses to the NLB for each Availability Zone. Create a target group and register the EC2 instances with the NLB. Create a new A (alias) record set named my.service.com, and assign the NLB DNS name to the record set.D. Create an Amazon ECS cluster and a service definition for the application. Create and assign public IP address for each host in the cluster. Create an Application Load Balancer (ALB) and expose the static TCP port. Create a target group and assign the ECS service definition name to the ALB. Create a new CNAME record set and associate the public IP addresses to the record set. Provide the Elastic IP addresses of the Amazon EC2 instances to the other companies to add to their allow lists.Answer: CExplanation:NLB with one Elastic IP per AZ to handle TCP traffic. Alias record set named my.service.com.
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-to-elb-load-balancer.html
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-to-elb-load-balancer.htmlQUESTION 176A company is running multiple workloads in the AWS Cloud. The company has separate units for software development. The company uses AWS Organizations and federation with SAML to give permissions to developers lo manage resources m their AWS accounts. The development units each deploy their production workloads into a common production account.Recently, an incident occurred in the production account in which members of a development unit terminated an EC2 instance that belonged to a different development unit.A solutions architect must create u solution that prevents a similar incident from happening in the future.The solution also must allow developers the possibility to manage the instances used for their workloads.Which strategy will meet these requirements?A. Create separate OUs in AWS Organizations for each development unit.Assign the created OUs to the company AWS accounts. Create separate SCPs with a deny action and a StringNotEquals condition for the DevelopmentUnit resource tag that matches the development unit name.Assign the SCP to the corresponding OU.B. Pass an attribute for DevelopmentUnit as an AWS Security Token Service (AWS STS) session tag during SAML federation.Update the AM policy for the developers'assumed IAM role with a deny action and a StringNotEquals condition for the DevelopmentUnit resource lag and aws:PrincipalTag/'DevelopmentUnit.C. Pass an attribute for DevelopmentUnit as an AWS Security Token Service (AWS STS) session tag curing SAML federation.Create an SCP with an allow action and a StringEquals condition for the DevelopmentUnit resource tag and aws:PrincipalTag/DevelopmentUnit.Assign the SCP to the root OU.D. Create separate IAM policies for each development unit. For every IAM policy, add an allow action and a StringEquals condition for the DevelopmentUnit resource tag and the development unit name.During SAML federation, use AWS Security Token Service (AWS STS) to assign the IAN' policy and match the development unit name to the assumed IAM role.Answer: BExplanation:A - Does not make much sense. An account can only belong to one OU. This is a single production account so it can't be in multiple OUs.B - Session tag is used to identify which business unit a user is part of. IAM policy prevent them from modifying resources for any business unit but their own.C - This does not restrict any existing permissions so users can still modify resources from different business units.D - STS cannot be used to assign a policy to an IAM role. A policy has to be assigned to the role before authentication occurs.
https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_abac-saml.htmlQUESTION 177A Solutions Architect is working with a company that is extremely sensitive to its IT costs and wishes to implement controls that will result in a predictable AWS spend each month.Which combination of steps can help the company control and monitor its monthly AWS usage to achieve a cost that is as close as possible to the target amount? (Choose three.)A. Implement an IAM policy that requires users to specify a `workload'

tag for cost allocation when launching Amazon EC2 instances.B. Contact AWS Support and ask that they apply limits to the account so that users are not able to launch more than a certain number of instance types.C. Purchase all upfront Reserved Instances that cover 100% of the account's expected Amazon EC2 usage.D. Place conditions in the users' IAM policies that limit the number of instances they are able to launch.E. Define `workload' as a cost allocation tag in the AWS Billing and Cost Management console.F. Set up AWS Budgets to alert and notify when a given workload is expected to exceed a defined cost.Answer: AEFExplanation:B: not feasible.C: Not everything is applicable to RI. E.g. S3 does not have RI.D: If they chose a very big instance, the bill could still be big.QUESTION 178A company wants to migrate its website from an on-premises data center onto AWS. At the same time, it wants to migrate the website to a containerized microservice-based architecture to improve the availability and cost efficiency. The company's security policy states that privileges and network permissions must be configured according to best practice, using least privilege.A Solutions Architect must create a containerized architecture that meets the security requirements and has deployed the application to an Amazon ECS cluster.What steps are required after the deployment to meet the requirements? (Choose two.)A. Create tasks using the bridge network mode.B. Create tasks using the awsvpc network mode.C. Apply security groups to Amazon EC2 instances, and use IAM roles for EC2 instances to access other resources.D. Apply security groups to the tasks, and pass IAM credentials into the container at launch time to access other resources.E. Apply security groups to the tasks, and use IAM roles for tasks to access other resources.Answer: BEExplanation:
https://aws.amazon.com/about-aws/whats-new/2017/11/amazon-ecs-introduces-awsvpc-networking-mode-for-containers-to-support-full-networking-capabilities/QUESTION 179A company has implemented a global multiplayer gaming platform. The platform requires gaming clients to have reliable, low-latency access to the server infrastructure that is hosted on a fleet of Amazon EC2 instances in a single AWS RegionThe gaming clients use a custom TCP protocol to connect to the server infrastructure.The application architecture requires client IP addresses to be available to the server software.Which solution meets these requirements?A. Create a Network Load Balancer (NLB), and add the EC2 instances to a target groupCreate an Amazon CloudFront Real Time Messaging Protocol (RTMP) distribution and configure the origin to point to the DNS endpoint of the NLB Use proxy protocol version 2 headers to preserve client IP addressesB. Use an AWS Direct Connect gateway to connect multiple Direct Connect locations in different Regions globallyConfigure Amazon Route 53 with geolocation routing to send traffic to the nearest Direct Connect locationAssociate the VPC that contains the EC2 instances with the Direct Connect gatewayC. Create an accelerator in AWS Global Accelerator and configure the listener to point to a single endpoint groupAdd each of the EC2 instances as endpoints to the endpoint groupConfigure the endpoint group weighting equally across all of the EC2 endpointsD. Create an Application Load Balancer (ALB) and add the EC2 instances to a target groupCreate a set of Amazon Route 53 latency-based alias records that point to the DNS endpoint of the ALBUse X-Forwarded-For headers to preserve client IP addressesAnswer: AQUESTION 180A solutions architect is importing a VM from an on-premises environment by using the Amazon EC2 VM Import feature of AWS Import/Export. The solutions architect has created an AMI and has provisioned an Amazon EC2 instance that is based on that AMI. The EC2 instance runs inside a public subnet in a VPC and has a public IP address assigned. The EC2 instance does not appear as a managed instance in the AWS Systems Manager console. Which combination of steps should the solutions architect take to troubleshoot this issue? (Choose two.)A. Verify that Systems Manager Agent is installed on the instance and is runningB. Verify that the instance is assigned an appropriate IAM role for Systems ManagerC. Verify the existence of a VPC endpoint on the VPCD. Verify that the AWS Application Discovery Agent is configuredE. Verify the correct configuration of service-linked roles for Systems ManagerAnswer: ABExplanation:
https://aws.amazon.com/premiumsupport/knowledge-center/systems-manager-ec2-instance-not-appear/QUESTION 181A retail company has a small ecommerce web application that uses an Amazon RDS for PostgreSQL DB instance. The DB instance is deployed with the Multi-AZ option turned on. Application usage recently increased exponentially and users experienced frequent http 503 errors. Users reported the errors, and the company's reputation suffered. The company could not identify a definitive root cause.The company wants to improve its operational readiness and receive alerts before users notice an incident.The company also wants to collect enough information to determine the root cause of any future incident.Which solution will meet these requirements with the LEAST operational overhead?A. Turn on Enhanced Monitoring for the DB instanceModify the corresponding parameter group to turn on query logging for all the slow queriesCreate Amazon CloudWatch alarmsSet the alarms to appropriate thresholds that are based on performance metrics in CloudWatchB. Turn on Enhanced Monitoring and Performance Insights for the DB instanceCreate Amazon CloudWatch alarmsSet the alarms to appropriate thresholds that are based on performance metrics in CloudWatchC. Turn on log exports to Amazon CloudWatch for the PostgreSQL logs on the DB instanceAnalyze the logs by using Amazon Elasticsearch Service (Amazon ES) and KibanaCreate a dashboard in KibanaConfigure alerts that are based on the metrics that are collectedD. Turn on Performance Insights for the DB instanceModify the corresponding parameter group to turn on query

logging for all the slow queriesCreate Amazon CloudWatch alarmsSet the alarms to appropriate thresholds that are based on performance metrics in CloudWatchAnswer: BExplanation:The cause of the issue is not known (i.e. it might not be slow queries) and RDS has SQL statistics in Performance Insight to investigate.
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/sql-statistics.html
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PerfInsights.Cloudwatch.html
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PerfInsights_Counters.htmlQUESTION 182A company has developed a new release of a popular video game and wants to make it available for public download. The new release package is approximately 5 GB in size. The company provides downloads for existing releases from a Linux-based, publicly facing FTP site hosted in an on-premises data center. The company expects the new release will be downloaded by users worldwide. The company wants a solution that provides improved download performance and low transfer costs, regardless of a user's location.Which solutions will meet these requirements?A. Store the game files on Amazon EBS volumes mounted on Amazon EC2 instances within an Auto Scaling group. Configure an FTP service on the EC2 instances. Use an Application Load Balancer in front of the Auto Scaling group. Publish the game download URL for users to download the package.B. Store the game files on Amazon EFS volumes that are attached to Amazon EC2 instances within an Auto Scaling group. Configure an FTP service on each of the EC2 instances. Use an Application Load Balancer in front of the Auto Scaling group. Publish the game download URL for users to download the package.C. Configure Amazon Route 53 and an Amazon S3 bucket for website hosting. Upload the game files to the S3 bucket. Use Amazon CloudFront for the website. Publish the game download URL for users to download the package.D. Configure Amazon Route 53 and an Amazon S3 bucket for website hosting. Upload the game files to the S3 bucket. Set Requester Pays for the S3 bucket. Publish the game download URL for users to download the package.Answer: CExplanation:
https://aws.amazon.com/blogs/aws/amazon-cloudfront-support-for-20-gb-objects/QUESTION 183A finance company is storing financial records in an Amazon S3 bucket. The company persists a record for every financial transaction.According to regulatory requirements, the records cannot be modified for at least 1 year after they are written.The records are read on a regular basis and must be immediately accessible.Which solution will meet these requirements?A. Create a new S3 bucket.Turn on S3 Object Lock, set a default retention period of 1 year, and set the retention mode to compliance mode.Store all records in the new S3 bucket.B. Create an S3 Lifecycle rule to immediately transfer new objects to the S3 Glacier storage tierCreate an S3 Glacier Vault Lock policy that has a retention period of 1 year.C. Create an S3 Lifecycle rule to immediately transfer new objects to the S3 Intelligent-Tiering storage tier.Set a retention period of 1 year.D. Create an S3 bucket policy with a Deny action for PutObject operations with a condition where the s3:x-amz-object-retention header is not equal to 1 year.Answer: AQUESTION 184A digital marketing company has multiple AWS accounts that belong to various teams. The creative team uses an Amazon S3 bucket in its AWS account to securely store images and media files that are used as content for the company's marketing campaigns. The creative team wants to share the S3 bucket with the strategy team so that the strategy team can view the objects.A solutions architect has created an IAM role that is named strategy_reviewer in the Strategy account. The solutions architect also has set up a custom AWS Key Management Service (AWS KMS) key in the Creative account and has associated the key with the S3 bucket. However, when users from the Strategy account assume the IAM role and try to access objects in the S3 bucket, they receive an Account. The solutions architect must ensure that users in the Strategy account can access the S3 bucket. The solution must provide these users with only the minimum permissions that they need.Which combination of steps should the solutions architect take to meet these requirements? (Choose three.)A. Create a bucket policy that includes read permissions for the S3 bucket.Set the principal of the bucket policy to the account ID of the Strategy accountB. Update the strategy_reviewer IAM role to grant full permissions for the S3 bucket and to grant decrypt permissions for the custom KMS key.C. Update the custom KMS key policy in the Creative account to grant decrypt permissions to the strategy_reviewer IAM role.D. Create a bucket policy that includes read permissions for the S3 bucket.Set the principal of the bucket policy to an anonymous user.E. Update the custom KMS key policy in the Creative account to grant encrypt permissions to the strategy_reviewer IAM role.F. Update the strategy_reviewer IAM role to grant read permissions for the S3 bucket and to grant decrypt permissions for the custom KMS keyAnswer: ACFExplanation:
https://aws.amazon.com/premiumsupport/knowledge-center/cross-account-access-denied-error-s3/In addition to the url above, you can eliminate the 3 of the answers easily:B- wrong becuase of the "full access"D- wrong because of the "anonymous user"E- wrong because of the "encrypt" - u need decrypt permissionQUESTION 185A company wants to deploy an API to AWS. The company plans to run the API on AWS Fargate behind a load balancer. The API requires the use of header-based routing and must be accessible from on- premises networks through an AWS Direct Connect connection and a private VIF. The company needs to add the client IP addresses that connect to the API to an allow list in AWS. The company also needs to add the IP addresses of the API to the allow list. The company's security team will allow /27 CIDR ranges to be added to the allow list. The solution must minimize

complexity and operational overhead.Which solution will meet these requirements?A. Create a new Network Load Balancer (NLB) in the same subnets as the Fargate task deployments.Create a security group that includes only the client IP addresses that need access to the API.Attach the new security group to the Fargate tasks.Provide the security team with the NLB's IP addresses for the allow list.B. Create two new /27 subnets.Create a new Application Load Balancer (ALB) that extends across the new subnets. Create a security group that includes only the client IP addresses that need access to the API.Attach the security group to the ALB. Provide the security team with the new subnet IP ranges for the allow list.C. Create two new /27 subnets.Create a new Network Load Balancer (NLB) that extends across the new subnets.Create a new Application Load Balancer (ALB) within the new subnets. Create a security group that includes only the client IP addresses that need access to the API. Attach thesecurity group to the ALB. Add the ALB's IP addresses as targets behind the NLB.Provide the security team with the NLB's IP addresses for the allow list.D. Create a new Application Load Balancer (ALB) in the same subnets as the Fargate task deployments.Create a security group that includes only the client IP addresses that need access to the API.Attach the security group to the ALB.Provide the security team with the ALB's IP addresses for the allow list.Answer: BExplanation:Since the security group will permit /27 CIDR ranges to be added to the allow list, we do not need to know what the actual IPs are of the ALBs (as they are dynamic). ALB is required as it operates at L7, needed for head-based routing.QUESTION 186A company is running an application in the AWS Cloud. The application runs on containers in an Amazon Elastic Container Service (Amazon ECS) cluster. The ECS tasks use the Fargate launch type. The application's data is relational and is stored in Amazon Aurora MySQL. To meet regulatory requirements, the application must be able to recover to a separate AWS Region in the event of an application failure. In case of a failure, no data can be lost.Which solution will meet these requirements with the LEAST amount of operational overhead?A. Provision an Aurora Replica in a different Region.B. Set up AWS DataSync for continuous replication of the data to a different Region.C. Set up AWS Database Migration Service (AWS DMS) to perform a continuous replication of the data to a different Region.D. Use Amazon Data Lifecycle Manager {Amazon DLM) to schedule a snapshot every 5 minutes.Answer: AExplanation:When you provision an Aurora Replica in a different AWS Region, the replica is kept in sync with the primary database using Aurora's replication capabilities. In the event of a failure in the primary Region, you can promote the Aurora Replica to become the new primary database, which allows you to continue operations with no data loss.However, provisioning and maintaining an Aurora Replica in a different AWS Region requires ongoing management and monitoring to ensure that it stays in sync with the primary databaseQUESTION 187A software company has deployed an application that consumes a REST API by using Amazon API Gateway. AWS Lambda functions, and an Amazon DynamoDB table. The application is showing an increase in the number of errors during PUT requests. Most of the PUT calls come from a small number of clients that are authenticated with specific API keys. A solutions architect has identified that a large number of the PUT requests originate from one client. The API is noncritical, and clients can tolerate retries of unsuccessful calls. However, the errors are displayed to customers and are causing damage to the API's reputation.What should the solutions architect recommend to improve the customer experience?A. Implement retry logic with exponential backoff and irregular variation in the client application.Ensure that the errors are caught and handled with descriptive error messages.B. Implement API throttling through a usage plan at the API Gateway level.Ensure that the client application handles code 429 replies without error.C. Turn on API caching to enhance responsiveness for the production stage.Run 10-minute load tests. Verify that the cache capacity is appropriate for the workload.D. Implement reserved concurrency at the Lambda function level to provide the resources that are needed during sudden increases in traffic.Answer: AExplanation:Custermer can accept delay /even failed attempts. https://aws.amazon.com/premiumsupport/knowledge-center/aws-batch-requests-error/QUESTION 188A company is planning to migrate its on-premises data analysis application to AWS. The application is hosted across a fleet of servers and requires consistent system time. The company has established an AWS Direct Connect connection from its on-premises data center to AWS. The company has a high-precision stratum-0 atomic dock network appliance that acts as an NTP source for all on-premises servers.After the migration to AWS is complete, the clock on all Amazon EC2 instances that host the application must be synchronized with the on-premises atomic clock network appliance.Which solution will meet these requirements with the LEAST administrative overhead?A. Configure a DHCP options set with the on-premises NTP server address.Assign the options set to the VPC.Ensure that NTP traffic is allowed between AWS and the on-premises networks.B. Create a custom AMI to use the Amazon Time Sync Service at 169.254.169.123Use this AMI for the applicationUse AWS Config to audit the NTP configuration.C. Deploy a third-party time server from the AWS Marketplace.Configure the time server to synchronize with the on-premises atomic clock network appliance.Ensure that NTP traffic is allowed inbound in the network ACLs for the VPC that contains the third-party server. D. Create an IPsec VPN tunnel from the on-premises atomic clock network appliance to the VPC to encrypt the traffic over the Direct Connect connection.Configure the VPC route tables to direct NTP traffic over the tunnel.Answer: AExplanation:Setting a VPC DHCP options set with your on-prem NTP server IP makes every EC2 instance automatically use that clock source. Traffic

rides the existing Direct Connect, so no extra tunnels or servers are needed ? minimal admin overhead while meeting the requirement to sync to the on-prem atomic clock.QUESTION 189A company has an on-premises Microsoft SQL Server database that writes a nightly 200 GB export to a local drive. The company wants to move the backups to more robust cloud storage on Amazon S3. The company has set up a 10 Gbps AWS Direct Connect connection between the on-premises data center and AWS. Which solution meets these requirements Most cost effectively?A. Create a new S3 bucket Deploy an AWS Storage Gateway file gateway within the VPC that is connected to the Direct Connect connection. Create a new SMB file share. Write nightly database exports to the new SMB file share.B. Create an Amzon FSx for Windows File Server Single-AZ file system within the VPC that is connected to the Direct Connect connection. Create a new SMB file share. Write nightly database exports to an SMB file share on the Amazon FSx file system Enable backups.C. Create an Amazon FSx for Windows File Server Multi-AZ system within the VPC that is connected to the Direct Connect connection. Create a new SMB file share. Write nightly database exports to an SMB file share on the Amazon FSx file system. Enable nightly backups.D. Create a new S3 buckets. Deploy an AWS Storage Gateway volume gateway within the VPC that is connected to the Direct Connect connection. Create a new SMB file share. Write nightly database exports to the new SMB file share on the volume gateway, and automate copies of this data to an S3 bucket.Answer: A Explanation:https://aws.amazon.com/storagegateway/file/?nc=sn&loc=2&dn=2QUESTION 190A company is processing videos in the AWS Cloud by using Amazon EC2 instances in an Auto Scaling group. It takes 30 minutes to process a video. Several EC2 instances scale in and out depending on the number of videos in an Amazon Simple Queue Service (Amazon SQS) queue. The company has configured the SQS queue with a redrive policy that specifies a target dead-letter queue and a maxReceiveCount of 1. The company has set the visibility timeout for the SQS queue to 1 hour. The company has set up an Amazon CloudWatch alarm to notify the development team when there are messages in the dead-letter queue.Several times during the day, the development team receives notification that messages are in the dead-letter queue and that videos have not been processed properly. An investigation finds no errors in the application logs.How can the company solve this problem?A. Turn on termination protection for the EC2 instances.B. Update the visibility timeout for the SOS queue to 3 hours.C. Configure scale-in protection for the instances during processing.D. Update the redrive policy and set maxReceiveCount to 0.Answer: CQUESTION 191An ecommerce company runs its infrastructure on AWS. The company exposes its APIs to its web and mobile clients through an Application Load Balancer (ALB) in front of an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The EKS cluster runs thousands of pods that provide the APIs.After extending delivery to a new continent, the company adds an Amazon CloudFront distribution and sets the ALB as the origin. The company also adds AWS WAF to its architecture.After implementation of the new architecture, API calls are significantly. However, there is a sudden increase in http status code 504 (Gateway Timeout) errors and http status code 502 (Bad Gateway) errors. This increase in errors seems to be for a specific domain.Which factors could be a cause of these errors? (Choose two.)A. AWS WAF is blocking suspicious requests.B. The origin is not properly configured in CloudFront.C. There is an SSL/TLS handshake issue between CloudFront and the origin.D. EKS Kubernetes pods are being cycled.E. Some pods are taking more than 30 seconds to answer API calls.Answer: AEQUESTION 192A gaming company created a game leaderboard by using a Multi-AZ deployment of an Amazon RDS database. The number of users is growing, and the queries to get individual player rankings are getting slower over time. The company expects a surge in users for an upcoming version and wants to optimize the design for scalability and performance.Which solution will meet these requirements?A. Migrate the database to Amazon DynamoDB. Store the leader different tables.Use Apache HiveQL JOIN statements to build the leaderboardB. Keep the leaderboard data in the RDS DB instance.Provision a Multi-AZ deployment of an Amazon ElastiCache for Redis cluster.C. Stream the leaderboard data by using Amazon Kinesis Data Firehose with an Amazon S3 bucket as the destination.Query the S3 bucket by using Amazon Athena for the leaderboard.D. Add a read-only replica to the RDS DB instance. Add an RDS Proxy database proxy.Answer: DExplanation:RDS Proxy makes applications more resilient to database failures by automatically connecting to a standby DB instance while preserving application connections. Using RDS Proxy, you can handle unpredictable surges in database traffic that otherwise might cause issues due to oversubscribing connections or creating new connections at a fast rate. https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/rds-proxy.htmlQUESTION 193A company has a serverless multi-tenant content management system on AWS. The architecture contains a web-based front end that interacts with an Amazon API Gateway API that uses a custom AWS Lambda authorizer. The authorizer authenticates a user to its tenant ID and encodes the information in a JSON Web Token (JWT) token. After authentication, each API call through API Gateway targets a Lambda function that interacts with a single Amazon DynamoDB table to fulfill requests.To comply with security standards, the company needs a stronger isolation between tenants. The company will have hundreds of customers within the first year.Which solution will meet these requirements with the LEAST operational overhead?A. Create a DynamoDB table for each tenant by using the tenant ID in the table name.Create a service that uses the JWT token to retrieve the appropriate Lambda execution role that is

tenant-specific.Attach IAM policies to the execution role to allow access only to the DynamoDB table for the tenant.B. Add tenant ID information to the partition key of the DynamoDB table.Create a service that uses the JWT token to retrieve the appropriate Lambda execution role that is tenant-specific.Attach IAM policies to the execution role to allow access to items in the table only when the key matches the tenant ID.C. Create a separate AWS account for each tenant of the application.Use dedicated infrastructure for each tenant.Ensure that no cross-account network connectivity exists.D. Add tenant ID as a sort key in every DynamoDB table.Add logic to each Lambda function to use the tenant ID that comes from the JWT token as the sort key in every operation on the DynamoDB table.Answer: BExplanation:Rather than creating table for each tenant, its better to use partition key in the already available table. This can be achieved with the LEAST operational.
https://aws.amazon.com/blogs/apn/multi-tenant-storage-with-amazon-dynamodb/QUESTION 194A 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 to migrate the data 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. 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.Answer: DExplanation: Create an AWS Snowball import job. Configure a server in the company TMs 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/QUESTION 195A greeting card company recently advertised that customers could send cards to their favorite celebrities through the company's platform. Since the advertisement was published, the platform has received constant traffic from 10,000 unique users each second.The platform runs on m5.xlarge Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an Auto Scaling group and use a custom AMI that is based on Amazon Linux. The platform uses a highly available Amazon Aurora MySQL DB cluster that uses primary and reader endpoints. The platform also uses an Amazon ElastiCache for Redis cluster that uses its cluster endpoint.The platform generates a new process for each customer and holds open database connections to MySQL for the duration of each customer's session. However, resource usage for the platform is low. Many customers are reporting errors when they connect to the platform. Logs show that connections to the Aurora database are failing. Amazon CloudWatch metrics show that the CPU load is

low across the platform and that connections to the platform are successful through the ALB.Which solution will remediate the errors MOST cost-effectively?A. Set up an Amazon CloudFront distribution. Set the ALB as the origin. Move all customer traffic to the CloudFront distribution endpoint.B. Use Amazon RDS Proxy. Reconfigure the database connections to use the proxy.C. Increase the number of reader nodes in the Aurora MySQL cluster.D. Increase the number of nodes in the ElastiCache for Redis cluster.Answer: BExplanation:https://aws.amazon.com/rds/proxy/faqs/QUESTION 196A company runs a highly available data collection application on Amazon EC2 in the eu-north-1 Region. The application collects data from end-user devices and writes records to an Amazon Kinesis data stream and a set of AWS Lambda functions that process the records .The company persists the output of the record processing to an Amazon S3 bucket in eu-north-1. The company uses the data in the S3 bucket as a data source for Amazon Athena.A. In each of the Iwo new Regions set up the Lambda functions to run in a VPC. Set up an S3 gateway endpoint in that VPC.B. Turn on S3 Transfer Acceleration on the S3 bucket in eu-north-1. Change the application to use the new S3 accelerated endpoint when the application uploads data to the S3 bucket.C. Create an S3 bucket in each of the two new Regions. Set the application in each new Region to upload to its respective S3 bucket. Set up S3 Cross-Region Replication to replicate data to the S3 bucket in eu-north-1.D. Increase the memory requirements of the Lambda functions to ensure that they have multiple cores available. Use the multipart upload feature when the application uploads data to Amazon S3 Lambda.Answer: AQUESTION 197A company wants to use Amazon S3 to back up its on-premises file storage solution. The company's on-premises file storage solution supports NFS, and the company wants its new solution to support NFS. The company wants to archive the backup files after 5 days. If the company needs archived files for disaster recovery, the company is willing to wait a few days for the retrieval of those files.Which solution meets these requirements MOST cost-effectively?A. Deploy an AWS Storage Gateway files gateway that is associated with an S3 bucket. Move the files from the on-premises file storage solution to the file gateway. Create an S3 Lifecycle rule to move the file to S3 Standard-Infrequent Access (S3 Standard-IA) after 5 days.B. Deploy an AWS Storage Gateway volume gateway that is associated with an S3 bucket. Move the files from the on-premises file storage solution to the volume gateway. Create an S3 Lifecycle rule to move the files to S3 Glacier Deep Archive after 5 days.C. Deploy an AWS Storage Gateway tape gateway that is associated with an S3 bucket. Move the files from the on-premises file storage solution to the tape gateway. Create an S3 Lifecycle rule to move the files to S3 Standard-Infrequent Access (S3 Standard-IA) after 5 days.D. Deploy an AWS Storage Gateway file gateway that is associated with an S3 bucket. Move the files from the on-premises file storage solution to the tape gateway. Create an S3 Lifecycle rule to move the files to S3 Standard-Infrequent Access (S3 Standard-IA) after 5 days.E. Deploy an AWS Storage Gateway file gateway that is associated with an S3 bucket. Move the files from the on-premises file storage solution to the file gateway. Create an S3 Lifecycle rule to move the files to S3 Glacier Deep Archive after 5 days.Answer: EExplanation:File gateway support NFS protocol, while volume gateway support iCSI protocol. And we need glacier deep archive to save cost, cause the company willing to wait for few days retrival time.QUESTION 198A large company runs workloads in VPCs that are deployed across hundreds of AWS accounts. Each VPC consists of public subnets and private subnets that span across multiple Availability Zones. NAT gateways are deployed in the public subnets and allow outbound connectivity to the internet from the private subnets. A solutions architect is working on a hub-and-spoke design. All private subnets in the spoke VPCs must route traffic to the internet through an egress VPC. The solutions architect already has deployed a NAT gateway in an egress VPC in a central AWS account.Which set of additional steps should the solutions architect take to meet these requirements?A. Create peering connections between the egress VPC and the spoke VPCs. Configure the required routing to allow access to the internet.B. Create a transit gateway, and share it with the existing AWS accounts. Attach existing VPCs to the transit gateway. Configure the required routing to allow access to the internet.C. Create a transit gateway in every account. Attach the NAT gateway to the transit gateways. Configure the required routing to allow access to the internet.D. Create an AWS PrivateLink connection between the egress VPC and the spoke VPCs. Configure the required routing to allow access to the internet.Answer: B Explanation:
https://docs.aws.amazon.com/whitepapers/latest/building-scalable-secure-multi-vpc-network-infrastructure/centralized-egress-to-internet.htmlay.htmlQUESTION 199A company has an application. Once a month, the application creates a compressed file that contains every object within an Amazon S3 bucket. The total size of the objects before compression is 1 TB. The application runs by using a scheduled cron job on an Amazon EC2 instance that has a 5 TB Amazon Elastic Block Store (Amazon EBS) volume attached. The application downloads all the files from the source S3 bucket to the EBS volume, compresses the file, and uploads the file to a target S3 bucket. Every invocation of the application takes 2 hours from start to finish.Which combination of actions should a solutions architect take to OPTIMIZE costs for this application? (Choose two.)A. Migrate the application to run an AWS Lambda function. Use Amazon EventBridge (Amazon CloudWatch Events) to schedule the Lambda function to run once each month.B. Configure the application to download the source files by using streams. Direct the streams into a compression library.

Direct the output of the compression library into a target object in Amazon S3.C. Configure the application to download the source files from Amazon S3 and save the files to local storage. Compress the files and upload them to Amazon S3.D. Configure the application to run as a container in AWS Fargate. Use Amazon EventBridge (Amazon CloudWatch Events) to schedule the task to run once each month.E. Provision an Amazon Elastic File System (Amazon EFS) file system. Attach the file system to the AWS Lambda function.Answer: BDExplanation:You can write S3 streams in mutliple languages using its SDK, there is no need to download the files.You should use Fargate and not Lambda becuase the proccesing time is bigger than 15 minutes. In Fargate you only pay for the resources you consume while your container is running instead of running the EC2 instance the all month.QUESTION 200A company is launching a web-based application in multiple regions around the world. The application consists of both static content stored in a private Amazon S3 bucket and dynamic content hosted in Amazon ECS containers content behind an Application Load Balancer (ALB). The company requires that the static and dynamic application content be accessible through Amazon CloudFront only.Which combination of steps should a solutions architect recommend to restrict direct content access to CloudFront? (Choose three.)A. Create a web ACL in AWS WAF with a rule to validate the presence of a custom header and associate the web ACL with the ALB.B. Create a web ACL in AWS WAF with a rule to validate the presence of a custom header and associate the web ACL with the CloudFront distribution.C. Configure CloudFront to add a custom header to origin requests.D. Configure the ALB to add a custom header to http requests.E. Update the S3 bucket ACL to allow access from the CloudFront distribution only.F. Create a CloudFront Origin Access Identity (OAI) and add it to the CloudFront distribution. Update the S3 bucket policy to allow access to the OAI only.Answer: ACFExplanation:Use CloudFront to add a custom header to all origin requests. Using AWS WAF, create a web rule that denies all requests without this custom header. Associate the web ACL to the CloudFront distribution is incorrect. If any new requests are going to CloudFront, they won't have the custom header initially so AWS WAF may block the request immediately. This could deny any new connections to CloudFront. Therefore, you need to associate the web ACL to the ALB, which is after the CloudFront adds the custom header.QUESTION 201A large company recently experienced an unexpected increase in Amazon RDS and Amazon DynamoDB costs. The company needs to increase visibility into delays of AWS Billing and Cost Management. There are various accounts associated with AWS Organizations, including many development and production accounts. There is no consistent tagging strategy across the organization, but there are guidelines in place that require all infrastructure to be deployed using AWS CloudFormation with consistent tagging. Management requires cost center numbers and project ID numbers for all existing and future DynamoDB tables and RDS distances.Which strategy should the solutions architect provide to meet these requirements?A. Use Tag Editor to tag existing resources.Create cost allocation tags to define the cost center and project ID and allow 24 hours for tags to propagate to existing resources.B. Use an AWS Config rule to alert the finance team of untagged resources.Create a centralized AWS Lambda based solution to tag untagged RDS databases and DynamoDB resources every hour using a cross-account role.C. Use Tag Editor to tag existing resources.Create cost allocation tags to define the cost center and project ID.Use SCPs to restrict resource creation that do not have the cost center and project ID on the resource.D. Create cost allocation tags to define the cost center and project ID and allow 24 hours for tags to propagate to existing resources.Update existing federated roles to restrict privileges to provision resources that do not include the cost center and project ID on the resource.Answer: CExplanation:For all present and prospective DynamoDB tables and RDS instances, management needs cost center and project ID numbers.QUESTION 202A company that has multiple AWS accounts is using AWS Organizations. The company's AWS accounts host VPCs, Amazon EC2 instances, and containers.The company's compliance team has deployed a security tool in each VPC where the company has deployments. The security tools run on EC2 instances and send information to the AWS account that is dedicated for the compliance team. The company has tagged all the compliance-related resources with a key of "costCenter" and a value or "compliance".The company wants to identify the cost of the security tools that are running on the EC2 instances so that the company can charge the compliance team's AWS account. The cost calculation must be as accurate as possible. What should a solutions architect do to meet these requirements?A. In the management account of the organization, activate the costCenter user-defined tag.Configure monthly AWS Cost and Usage Reports to save to an Amazon S3 bucket in the management account. Use the tag breakdown in the report to obtain the total cost for the costCenter tagged resources.B. In the member accounts of the organization, activate the costCenter user-defined tag. Configure monthly AWS Cost and Usage Reports to save to an Amazon S3 bucket in the management account.Schedule a monthly AWS Lambda function to retrieve the reports and calculate the total cost for the costCenter tagged resources.C. In the member accounts of the organization activate the costCenter user-defined tag. From the management account, schedule a monthly AWS Cost and Usage Report. Use the tag breakdown in the report to calculate the total cost for the costCenter tagged resources.D. Create a custom report in the organization view in AWS Trusted Advisor. Configure the report to generate a monthly billing summary for the costCenter tagged resources in the compliance team's AWS account.Answer: AExplanation:https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/custom-tags.html

https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/configurecostallocreport.htmlQUESTION 203A company has an organization that has many AWS accounts in AWS Organizations. A solutions architect must improve how the company manages common security group rules for the AWS accounts in the organization. The company has a common set of IP CIDR ranges in an allow list in each AWS account to allow access to and from the company's on-premises network. Developers within each account are responsible for adding new IP CIDR ranges to their security groups. The security team has its own AWS account. Currently, the security team notifies the owners of the other AWS accounts when changes are made to the allow list.The solutions architect must design a solution that distributes the common set of CIDR ranges across all accounts.Which solution meets these requirements with the LEAST amount of operational overhead?A. Set up an Amazon Simple Notification Service (Amazon SNS) topic in the security team's AWS account. Deploy an AWS Lambda function in each AWS account. Configure the Lambda function to run every time an SNS topic receives a message. Configure the Lambda function to take an IP address as input and add it to a list of security groups in the account. Instruct the security team to distribute changes by publishing messages to its SNS topic.B. Create new customer-managed prefix lists in each AWS account within the organization. Populate the prefix lists in each account with all internal CIDR ranges. Notify the owner of each AWS account to allow the new customer-managed prefix list IDs in their accounts in their security groups. Instruct the security team to share updates with each AWS account owner.C. Create a new customer-managed prefix list in the security team's AWS account. Populate the customer-managed prefix list with all internal CIDR ranges. Share the customer-managed prefix list with the organization by using AWS Resource Access Manager. Notify the owner of each AWS account to allow the new customer-managed prefix list ID in their security groups.D. Create an IAM role in each account in the organization. Grant permissions to update security groups. Deploy an AWS Lambda function in the security team's AWS account. Configure the Lambda function to take a list of internal IP addresses as input, assume a role in each organization account, and add the list of IP addresses to the security groups in each account.Answer: CExplanation: https://docs.aws.amazon.com/vpc/latest/userguide/managed-prefix-lists.htmlQUESTION 204A company uses AWS Organizations with a single OU named Production to manage multiple accounts. All accounts are members of the Production OU. Administrators use deny list SCPs in the root of the organization to manage access to restricted services. The company recently acquired a new business unit and invited the new unit's existing AWS account to the organization. Once onboarded, the administrators of the new business unit discovered that they are not able to update existing AWS Config rules to meet the company's policies.Which option will allow administrators to make changes and continue to enforce the current policies without introducing additional long-term maintenance?A. Remove the organization's root SCPs that limit access to AWS Config.Create AWS Service Catalog products for the company's standard AWS Config rules and deploy them throughout the organization, including the new account.B. Create a temporary OU named Onboarding for the new account.Apply an SCP to the Onboarding OU to allow AWS Config actions.Move the new account to the Production OU when adjustments to AWSConfig are complete.C. Convert the organization's root SCPs from deny list SCPs to allow list SCPs to allow the required services only.Temporally apply an SCP to the organization's root that allows AWS Config actions for principals only in the new account.D. Create a temporary OU named Onboarding for the new account. Apply an SCP to the Onboarding OU to allow AWS Config actions.Move the organization's root SCP to the Production OU.Move the new account to the Production OU when adjustments to AWS Config are complete.Answer: DExplanation:An SCP at a lower level can't add a permission after it is blocked by an SCP at a higher level. SCPs can only filter; they never add permissions. SO you need to create a new OU for the new account assign an SCP, and move the root SCP to Production OU. Then move the new account to production OU when AWS config is done.QUESTION 205A company is deploying a distributed in-memory database on a fleet of Amazon EC2 instances. The fleet consists of a primary node and eight worker nodes. The primary node is responsible for monitoring cluster health, accepting user requests, distributing user requests to worker nodes and sending an aggregate response back to a client. Worker nodes communicate with each other to replicate data partitions.The company requires the lowest possible networking latency to achieve maximum performance.Which solution will meet these requirements?A. Launch memory optimized EC2 instances in a partition placement groupB. Launch compute optimized EC2 instances in a partition placement groupC. Launch memory optimized EC2 instances in a cluster placement groupD. Launch compute optimized EC2 instances in a spread placement groupAnswer: CExplanation:The cluster placement group is indicated for high troughput and low network latency, while the partition placement group is used to avoid hw failures separating the instances in groups running over different hardware. https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html#placement-groups-clusterQUESTION 206A company is running an application in the AWS Cloud. The application consists of microservices that run on a fleet of Amazon EC2 instances in multiple Availability Zones behind an Application Load Balancer. The company recently added a new REST API that was implemented in Amazon API Gateway. Some of the older microservices that run on EC2 instances need to call this new API. The company does not want the API to be accessible from the public internet and does not want proprietary data to traverse the

public internet.What should a solutions architect do to meet these requirements?A. Create an AWS Site-to-Site VPN connection between the VPC and the API Gateway. Use API Gateway to generate a unique API key for each microservice. Configure the API methods to require the key.B. Create an interface VPC endpoint for API Gateway, and set an endpoint policy to only allow access to the specific API. Add a resource policy to API Gateway to only allow access from the VPC endpoint. Change the API Gateway endpoint type to private.C. Modify the API Gateway to use IAM authentication Update the IAM policy for the IAM role that is assigned to the EC2 instances to allow access to the API Gateway Move the API Gateway into a new VPC. Deploy a transit gateway and connect the VPCs.D. Create an accelerator in AWS Global Accelerator, and connect the accelerator to the API Gateway. Update the route table for all VPC subnets with a route to the created Global Accelerator endpoint IP address. Add an API key for each service to use for authentication.Answer: BExplanation:
https://docs.aws.amazon.com/apigateway/latest/developerguide/apigateway-vpc-endpoint-policies.htmlResources From:1.2025 Latest Braindump2go SAP-C02 Exam Dumps (PDF & VCE) Free Share:**https://www.braindump2go.com/sap-c02.html**2.2025 Latest Braindump2go SAP-C02 PDF and VCE Dumps Free Share:
https://drive.google.com/drive/folders/1cAO91F5KZT_krcyo83Yuh2sMJfyvKhZn?usp=sharing3.2025 Free Braindump2go SAP-C02Exam Questions Download:https://www.braindump2go.com/free-online-pdf/SAP-C02-VCE-Dumps(175-206).pdf Free Resources from Braindump2go,We Devoted to Helping You 100% Pass All Exams!