

[March-2022] 100% Exam Pass-SAP-C01 VCE Dumps Free from Braindump2go [Q983-Q1002]

March/2022 New Braindump2go SAP-C01 Exam Dumps with PDF and VCE Free Updated Today! Following are some new SAP-C01 Real Exam Questions! QUESTION 983 A company in the United States (US) has acquired a company in Europe. Both companies use the AWS Cloud. The US company has built a new application with a microservices architecture. The US company is hosting the application across five VPCs in the us-east-2 Region. The application must be able to access resources in one VPC in the eu-west-1 Region. However, the application must not be able to access any other VPCs. The VPCs in both Regions have no overlapping CIDR ranges. All Accounts are already consolidated in one organization in AWS Organizations. Which solution will meet these requirements MOST cost-effectively?

A. Create one transit gateway in eu-west-1. Attach the VPCs in us-east-2 and the VPC in eu-west-1 to the transit gateway. Create the necessary route entries in each VPC so that the traffic is routed through the transit gateway.

B. Create one transit gateway in each Region. Attach the involved subnets to the regional transit gateway. Create the necessary route entries in the associated route tables for each subnet so that the traffic is routed through the regional transit gateway. Peer the two transit gateways.

C. Create a full mesh VPC peering connection configuration between all the VPCs. Create the necessary route entries in each VPC so that the traffic is routed through the VPC peering connection.

D. Create one VPC peering connection for each VPC in us-east-2 to the VPC in eu-west-1. Create the necessary route entries in each VPC so that the traffic is routed through the VPC peering connection.

Answer: B

Explanation:

<https://docs.aws.amazon.com/vpc/latest/tgw/how-transit-gateways-work.html>

QUESTION 984 A United Kingdom (UK) company recently completed a successful proof of concept in Amazon WorkSpaces. The company also has a large office in the United States (US). Staff members from each office regularly travel between the two locations and need access to a corporate WorkSpace without any reconfiguration of their WorkSpaces client. The company has purchased a domain by using Amazon Route 53 for the connection alias. The company will use a Windows profile and document management solution. A solutions architect needs to design the full solution. The solution must use a configuration of WorkSpaces in two AWS Regions and must provide Regional resiliency. Which solution will meet these requirements?

A. Create a connection alias in a UK Region and a US Region. Associate the connection alias with a directory in the UK Region. Configure the DNS service for the domain in the connection alias. Configure a geolocation routing policy. Distribute the connection string to the WorkSpaces users.

B. Create a connection alias in a UK Region. Associate the connection alias with a directory in the UK Region. Configure the DNS service for the domain in the connection alias. Configure a weighted routing policy, with the UK Region set to 1 and a US Region set to 255. Distribute the connection string for the UK Region to the WorkSpaces users.

C. Create a connection alias in a UK Region and a US Region. Associate the connection aliases with a directory in each Region. Configure the DNS service for the domain in the connection alias. Configure a geolocation routing policy. Distribute the connection string to the WorkSpaces users.

D. Create a connection alias in a US Region. Associate the connection alias with a directory in the UK Region. Configure the DNS service for the domain in the connection alias. Configure a multivalue answer routing policy. Distribute the connection string for the US Region to the WorkSpaces users.

Answer: C

Explanation: <https://docs.aws.amazon.com/workspaces/latest/adminguide/cross-region-redirection.html>

QUESTION 985 A company is running a custom database in the AWS Cloud. The database uses Amazon EC2 for compute and uses Amazon Elastic Block Store (Amazon EBS) for storage. The database runs on the latest generation of EC2 instances and uses a General Purpose SSD (gp2) EBS volume for data. The current data volume has the following characteristics: The volume is 512 GB in size. The volume never goes above 256 GB utilization. The volume consistently uses around 1,500 IOPS. A solutions architect needs to conduct an analysis of the current database storage layer and make a recommendation about ways to reduce cost. Which solution will provide the MOST cost savings without impacting the performance of the database?

A. Convert the data volume to the Cloud HDD (sc1) type. Leave the volume as 512 GB. Set the volume IOPS to 1,500.

B. Convert the data volume to the Provisioned IOPS SSD (io2) type. Resize the volume to 256 GB. Set the volume IOPS to 1,500.

C. Convert the data volume to the Provisioned IOPS SSD (io2) Block Express type. Leave the volume as 512 GB. Set the volume IOPS to 1,500.

D. Convert the data volume to the General Purpose SSD (gp3) type. Resize the volume to 256 GB. Set the volume IOPS to 1,500.

Answer: C

QUESTION 986 A retail company is hosting an ecommerce website on AWS across multiple AWS Regions. The company wants the website to be operational at all times for online purchases. The website stores data in an Amazon RDS for MySQL DB instance. Which solution will provide the HIGHEST availability for the database?

A. Configure automated backups on Amazon RDS. In the case of disruption, promote an automated backup to be a standalone DB instance. Direct database traffic to the promoted DB instance. Create a replacement read replica that has the promoted DB instance as its source.

B. Configure global tables and read replicas on Amazon RDS. Activate the cross-Region scope. In the case of disruption, use AWS Lambda to copy the read replicas from one Region to another Region.

C.

Configure global tables and automated backups on Amazon RDS. In the case of disruption, use AWS Lambda to copy the read replicas from one Region to another Region.D. Configure read replicas on Amazon RDS. In the case of disruption, promote a cross-Region and read replica to be a standalone DB instance. Direct database traffic to the promoted DB instance. Create a replacement read replica that has the promoted DB instance as its source.Answer: AExplanation:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.htmlQUESTION 987A 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: AExplanation:

<https://aws.amazon.com/blogs/database/storing-sql-server-backups-in-amazon-s3-using-aws-storage-gateway/>QUESTION 988A video processing company wants to build a machine learning (ML) model by using 600 TB of compressed data that is stored as thousands of files in the company's on-premises network attached storage system. The company does not have the necessary compute resources on premises for ML experiments and wants to use AWS.The company needs to complete the data transfer to AWS within 3 weeks. The data transfer will be a one-time transfer. The data must be encrypted in transit. The measured upload speed of the company's internet connection is 100 Mbps, and multiple departments share the connection.Which solution will meet these requirements MOST cost-effectively?A. Order several AWS Snowball Edge Storage Optimized devices by using the AWS Management Console. Configure the devices with a destination S3 bucket. Copy the data to the devices. Ship the devices back to AWS.B. Set up a 10 Gbps AWS Direct Connect connection between the company location and the nearest AWS Region. Transfer the data over a VPN connection into the Region to store the data in Amazon S3.C. Create a VPN connection between the on-premises network storage and the nearest AWS Region. Transfer the data over the VPN connection.D. Deploy an AWS Storage Gateway file gateway on premises. Configure the file gateway with a destination S3 bucket. Copy the data to the file gateway.Answer: BExplanation:

<https://aws.amazon.com/blogs/networking-and-content-delivery/building-a-global-network-using-aws-transit-gateway-inter-region-peering/>QUESTION 989A company has several AWS accounts. A development team is building an automation framework for cloud governance and remediation processes. The automation framework uses AWS Lambda functions in a centralized account. A solutions architect must implement a least privilege permissions policy that allows the Lambda functions to run in each of the company's AWS accounts.Which combination of steps will meet these requirements? (Choose two.)A. In the centralized account, create an IAM role that has the Lambda service as a trusted entity. Add an inline policy to assume the roles of the other AWS accounts.B. In the other AWS accounts, create an IAM role that has minimal permissions. Add the centralized account's Lambda IAM role as a trusted entity.C. In the centralized account, create an IAM role that has roles of the other accounts as trusted entities. Provide minimal permissions.D. In the other AWS accounts, create an IAM role that has permissions to assume the role of the centralized account. Add the LambdaE. In the other AWS accounts, create an IAM role that has minimal permissions. Add the Lambda service as a trusted entity.Answer: ACEExplanation:

<https://aws.amazon.com/blogs/devops/how-to-centrally-manage-aws-config-rules-across-multiple-aws-accounts/>QUESTION 990A hedge fund company is developing a new web application to handle trades. Traders around the world will use the application. The application will handle hundreds of thousands of transactions, especially during overlapping work hours between Europe and the United States.According to the company's disaster recovery plan, the data that is generated must be replicated to a second AWS Region. Each transaction item will be less than 100 KB in size. The company wants to simplify the CI/CD pipeline as much as possible.Which combination of steps will meet these requirements MOST cost-effectively? (Choose two.)A. Deploy the application in multiple Regions. Use Amazon Route 53 latency-based routing to route users to the nearest deployment.B. Provision

an Amazon Aurora global database to persist data. Use Amazon ElastiCache to improve response time.C. Provision an Amazon CloudFront domain with the website as an origin. Restrict access to geographies where the usage is expected.D. Provision an Amazon DynamoDB global table. Use DynamoDB Accelerator (DAX) to improve response time.E. Provision an Amazon Aurora multi-master cluster to persist data. Use Amazon ElastiCache to improve response time.

Answer: ABExplanation:
<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html>QUESTION 991A company is migrating some of its applications to AWS. The company wants to migrate and modernize the applications quickly after it finalizes networking and security strategies. The company has set up an AWS Direct Connection connection in a central network account. The company expects to have hundreds of AWS accounts and VPCs in the near future. The corporate network must be able to access the resources on AWS seamlessly and also must be able to communicate with all the VPCs. The company also wants to route its cloud resources to the internet through its on-premises data center. Which combination of steps will meet these requirements? (Choose three.)

A. Create a Direct Connect gateway in the central account. In each of the accounts, create an association proposal by using the Direct Connect gateway and the account ID for every virtual private gateway.B. Create a Direct Connect gateway and a transit gateway in the central network account. Attach the transit gateway to the Direct Connect gateway by using a transit VIF.C. Provision an internet gateway. Attach the internet gateway to subnets. Allow internet traffic through the gateway.D. Share the transit gateway with other accounts. Attach VPCs to the transit gateway.E. Provision VPC peering as necessary.F. Provision only private subnets. Open the necessary route on the transit gateway and customer gateway to allow outbound internet traffic from AWS to flow through NAT services that run in the data center.

Answer: BDEExplanation:
<https://docs.aws.amazon.com/vpc/latest/tgw/tgw-dcg-attachments.html>QUESTION 992A company is building a

software-as-a-service (SaaS) solution on AWS. The company has deployed an Amazon API Gateway REST API with AWS Lambda integration in multiple AWS Regions and in the same production account. The company offers tiered pricing that gives customers the ability to pay for the capacity to make a certain number of API calls per second. The premium tier offers up to 3,000 calls per second, and customers are identified by a unique API key. Several premium tier customers in various Regions report that they receive error responses of 429 Too Many Requests from multiple API methods during peak usage hours. Logs indicate that the Lambda function is never invoked. What could be the cause of the error messages for these customers?A. The Lambda function reached its concurrency limit.B. The Lambda function reached its Region limit for concurrency.C. The company reached its API Gateway account limit for calls per second.D. The company reached its API Gateway default per-method limit for calls per second.

Answer: CQUESTION 993A company has registered 10 new domain names. The company uses the domains for online marketing. The company needs a solution that will redirect online visitors to a specific URL for each domain. All domains and target URLs are defined in a JSON document. All DNS records are managed by Amazon Route 53. A solutions architect must implement a redirect service that accepts HTTP and HTTPS requests. Which combination of steps should the solutions architect take to meet these requirements with the LEAST amount of operational effort? (Choose three.)

A. Create a dynamic webpage that runs on an Amazon EC2 instance. Configure the webpage to use the JSON document in combination with the event message to look up and respond with a redirect URL.B. Create an Application Load Balancer that includes HTTP and HTTPS listeners.C. Create an AWS Lambda function that uses the JSON document in combination with the event message to look up and respond with a redirect URL.

D. Use an Amazon API Gateway API with a custom domain to publish an AWS Lambda function.E. Create an Amazon CloudFront distribution. Deploy a Lambda@Edge function.F. Create an SSL certificate by using AWS Certificate Manager (ACM). Include the domains as Subject Alternative Names.

Answer: ABFExplanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/route-53-redirect-to-another-domain/>QUESTION 994A company asks a solution architect to optimize the cost of a solution. The solution handles requests from multiple customers. The solution includes a multi-tier architecture that uses Amazon API Gateway, AWS Lambda, AWS Fargate, Amazon Simple Queue Service (Amazon SQS), and Amazon EC2. In the current setup, requests go through API Gateway to Lambda and either start a container in Fargate or push a message to an SQS queue. An EC2 Fleet provides EC2 instances that serve as workers for the SQS queue. The EC2 Fleet scales based on the number of items in the SQS queue. Which combination of steps should the solutions architect recommend to reduce cost the MOST? (Choose three.)

A. Determine the minimum number of EC2 instances that are needed during a day. Reserve this number of instances in a 3-year plan with payment all upfront.B. Examine the last 6 months of compute utilization across the services. Use this information to determine the needed compute for the solution. Commit to a Savings Plan for this amount.C. Determine the average number of EC2 instances that are needed during a day. Reserve this number of instances in a 3-year plan with payment all upfront.D. Remove the SQS queue from the solution and from the solution infrastructure.E. Change the solution so that it runs as a container instead of on EC2 instances. Configure Lambda to start up the solution in Fargate by using environment variables to give the solution the message.F. Change the Lambda function so that it posts the message directly to the

EC2 instances through an Application Load Balancer. Answer: CDEExplanation:

<https://aws.amazon.com/ec2/pricing/reserved-instances/> QUESTION 995 A company is developing a messaging application that is based on a microservices architecture. A separate team develops each microservice by using Amazon Elastic Container Service (Amazon ECS). The teams deploy the microservices multiple times daily by using AWS CloudFormation and AWS CodePipeline. The application recently grew in size and complexity. Each service operates correctly on its own during development, but each service produces error messages when it has to interact with other services in production. A solutions architect must improve the application's availability. Which solution will meet these requirements with the LEAST amount of operational overhead? A. Add an extra stage to CodePipeline for each service. Use the extra stage to deploy each service to a test environment. Test each service after deployment to make sure that no error messages occur. B. Add an AWS::CodeDeployBlueGreen Transform section and Hook section to the template to enable blue/green deployments by using AWS CodeDeploy in CloudFormation. Configure the template to perform ECS blue/green deployments in production. C. Add an extra stage to CodePipeline for each service. Use the extra stage to deploy each service to a test environment. Write integration tests for each service. Run the tests automatically after deployment. D. Use an ECS DeploymentConfiguration parameter in the template to configure AWS CodeDeploy to perform a rolling update of the service. Use a CircuitBreaker property to roll back the deployment if any error occurs during deployment.

Answer: AExplanation: <https://aws.amazon.com/blogs/devops/using-aws-codepipeline-for-deploying-container-images-to-microservices-architecture-involving-aws-lambda-functions/> QUESTION 996 A company is migrating mobile banking applications to run on Amazon EC2 instances in a VPC. Backend service applications run in an on-premises data center. The data center has an AWS Direct Connect connection into AWS. The applications that run in the VPC need to resolve DNS requests to an on-premises Active Directory domain that runs in the data center. Which solution will meet these requirements with the LEAST administrative overhead? A. Provision a set of EC2 instances across two Availability Zones in the VPC as caching DNS servers to resolve DNS queries from the application servers within the VPC. B. Provision an Amazon Route 53 private hosted zone. Configure NS records that point to on-premises DNS servers. C. Create DNS endpoints by using Amazon Route 53 Resolver Add conditional forwarding rules to resolve DNS namespaces between the on-premises data center and the VPC. D. Provision a new Active Directory domain controller in the VPC with a bidirectional trust between this new domain and the on-premises Active Directory domain.

Answer: BExplanation: <https://aws.amazon.com/blogs/security/how-to-set-up-dns-resolution-between-on-premises-networks-and-aws-using-aws-directory-service-and-amazon-route-53/> QUESTION 997 A company has a new security policy. The policy requires the company to log any event that retrieves data from Amazon S3 buckets. The company must save these audit logs in a dedicated S3 bucket. The company created the audit logs S3 bucket in an AWS account that is designated for centralized logging. The S3 bucket has a bucket policy that allows write-only cross-account access. A solutions architect must ensure that all S3 object-level access is being logged for current S3 buckets and future S3 buckets. Which solution will meet these requirements? A. Enable server access logging for all current S3 buckets. Use the audit logs S3 bucket as a destination for audit logs. B. Enable replication between all current S3 buckets and the audit logs S3 bucket. Enable S3 Versioning in the audit logs S3 bucket. C. Configure S3 Event Notifications for all current S3 buckets to invoke an AWS Lambda function every time objects are accessed. Store Lambda logs in the audit logs S3 bucket. D. Enable AWS CloudTrail, and use the audit logs S3 bucket to store logs. Enable data event logging for S3 event sources, current S3 buckets, and future S3 buckets.

Answer: DExplanation: <https://docs.aws.amazon.com/awscloudtrail/latest/userguide/best-practices-security.html> QUESTION 998 A retail company is operating its ecommerce application on AWS. The application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The company uses an Amazon RDS DB instance as the database backend. Amazon CloudFront is configured with one origin that points to the ALB. Static content is cached. Amazon Route 53 is used to host all public zones. After an update of the application, the ALB occasionally returns a 502 status code (Bad Gateway) error. The root cause is malformed HTTP headers that are returned to the ALB. The webpage returns successfully when a solutions architect reloads the webpage immediately after the error occurs. While the company is working on the problem, the solutions architect needs to provide a custom error page instead of the standard ALB error page to visitors. Which combination of steps will meet this requirement with the LEAST amount of operational overhead? (Choose two.) A. Create an Amazon S3 bucket. Configure the S3 bucket to host a static webpage. Upload the custom error pages to Amazon S3. B. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function if the ALB health check response TargetFailedHealthChecks is greater than 0. Configure the Lambda function to modify the forwarding rule at the ALB to point to a public C. Modify the existing Amazon Route 53 records by adding health checks. Configure a fallback target if the health check fails. D. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function if the ALB health check response ElbInternalError is greater than 0. Configure the Lambda function to modify the forwarding rule at the ALB to point to a public accessible web server. E. Add a custom error response by configuring a CloudFront custom error page. Modify DNS records to

point to a publicly accessible web page. Answer: DEExplanation:

<https://aws.amazon.com/blogs/compute/scaling-amazon-ecs-services-automatically-using-amazon-cloudwatch-and-aws-lambda/>

QUESTION 999A 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: BExplanation: <https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html>

QUESTION 1000A financial services company sells its software-as-a-service (SaaS) platform for application compliance to large global banks. The SaaS platform runs on AWS and uses multiple AWS accounts that are managed in an organization in AWS Organizations. The SaaS platform uses many AWS resources globally. For regulatory compliance, all API calls to AWS resources must be audited, tracked for changes, and stored in a durable and secure data store. Which solution will meet these requirements with the LEAST operational overhead?A. Create a new AWS CloudTrail trail. Use an existing Amazon S3 bucket in the organization's management account to store the logs. Deploy the trail to all AWS Regions. Enable MFA delete and encryption on the S3 bucket.B. Create a new AWS CloudTrail trail in each member account of the organization. Create new Amazon S3 buckets to store the logs. Deploy the trail to all AWS Regions. Enable MFA delete and encryption on the S3 buckets.C. Create a new AWS CloudTrail trail in the organization's management account. Create a new Amazon S3 bucket with versioning turned on to store the logs. Deploy the trail for all accounts in the organization. Enable MFA delete and encryption on the S3 bucket.D. Create a new AWS CloudTrail trail in the organization's management account. Create a new Amazon S3 bucket to store the logs. Configure Amazon Simple Notification Service (Amazon SNS) to send log-file delivery notifications to an external management system that will track the logs. Enable MFA delete and encryption on the S3 bucket. Answer: DExplanation:

<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/creating-an-organizational-trail-in-the-console.html>

QUESTION 1001A retail company runs a business-critical web service on an Amazon Elastic Container Service (Amazon ECS) cluster that runs on Amazon EC2 instances. The web service receives POST requests from end users and writes data to a MySQL database that runs on a separate EC2 instance. The company needs to ensure that data loss does not occur. The current code deployment process includes manual updates of the ECS service. During a recent deployment, end users encountered intermittent 502 Bad Gateway errors in response to valid web requests. The company wants to implement a reliable solution to prevent this issue from recurring. The company also wants to automate code deployments. The solution must be highly available and must optimize cost-effectiveness. Which combination of steps will meet these requirements? (Choose three.)A. Run the web service on an ECS cluster that has a Fargate launch type. Use AWS CodePipeline and AWS CodeDeploy to perform a blue/green deployment with validation testing to update the ECS service.B. Migrate the MySQL database to run on an Amazon RDS for MySQL Multi-AZ DB instance that uses Provisioned IOPS SSD (io2) storage.C. Configure an Amazon Simple Queue Service (Amazon SQS) queue as an event source to receive the POST requests from the web service. Configure an AWS Lambda function to poll the queue. Write the data to the database.D. Run the web service on an ECS cluster that has a Fargate launch type. Use AWS CodePipeline and AWS CodeDeploy to perform a canary deployment to update the ECS service.E. Configure an Amazon Simple Queue Service (Amazon SQS) queue. Install the SQS agent on the containers that run in the ECS cluster to poll the queue. Write the data to the database.F. Migrate the MySQL database to run on an Amazon RDS for MySQL Multi-AZ DB instance that uses General Purpose SSD (gp3) storage. Answer: BCD

QUESTION 1002A company is using an Amazon EMR cluster to run its big data jobs. The cluster's jobs are invoked by AWS Step Functions Express Workflows that consume various Amazon Simple Queue Service (Amazon SQS) queues. The workload of this solution is variable and unpredictable. Amazon CloudWatch metrics show that the cluster's peak utilization is only 25% at times and that the cluster sits idle the rest of the time. A solutions architect must optimize the costs of the cluster without negatively impacting the time it takes to run the various jobs. What is the MOST cost-effective solution that meets these requirements?A. Modify the EMR cluster by turning on automatic scaling of the core nodes and task nodes with a custom policy that is based on cluster utilization. Purchase Reserved Instance capacity to cover the master node.B. Modify the EMR cluster to use an instance fleet of Dedicated On-Demand Instances for the master node and core nodes, and to use Spot Instances for the task

nodes. Define target capacity for each node type to cover the load.C. Purchase Reserved Instances for the master node and core nodes. Terminate all existing task nodes in the EMR cluster.D. Modify the EMR cluster to use capacity-optimized Spot Instances and a diversified task fleet. Define target capacity for each node type with a mix of On-Demand Instances and Spot Instances.

Answer: DExplanation:<https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-instance-fleet.html>Resources

From:1.2022 Latest Braindump2go SAP-C01 Exam Dumps (PDF & VCE) Free Share:

<https://www.braindump2go.com/aws-certified-solutions-architect-professional.html>2.2022 Latest Braindump2go SAP-C01 PDF and SAP-C01 VCE Dumps Free Share:

<https://drive.google.com/drive/folders/1wLkIVBV7ihIea0h2CrPoXpZliQHhVDh8?usp=sharing>3.2021 Free Braindump2go SAP-C01 Exam Questions Download:[https://www.braindump2go.com/free-online-pdf/SAP-C01-PDF-Dumps\(983-1002\).pdf](https://www.braindump2go.com/free-online-pdf/SAP-C01-PDF-Dumps(983-1002).pdf)Free Resources from Braindump2go, We Devoted to Helping You 100% Pass All Exams!