[2020-August-NewOffer SAP-C01 Dumps PDF SAP-C01 668Q for Free Downloading in Braindump2go(Q632-Q647)

August/2020 Latest Braindump2go SAP-C01 Exam Dumps with PDF and VCE Free Updated Today! Following are some new SAP-C01 Real Exam Questions!QUESTION 632A 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: BQUESTION 633A 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.r4.4xiarge 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 equivalently sized read replicas to provide reliability.D.   Reserve capacity for all EC2 instances and leverage Spot Instance pricing for the RDS database.Answer: BQUESTION 634During an audit, a security team discovered that a development team was putting IAM user secret access keys in their code and then committing it to an AWS CodeCommit repository. The security team wants to automatically find and remediate instances of this security vulnerability.Which solution will ensure that the credentials are appropriately secured automatically?A.   Run a script nightly using AWS Systems Manager.Run Command to search for credentials on the development instances. If found, use AWS Secrets Manager to rotate the credentials.B.   Use a scheduled AWS Lambda function to download and scan the application code from CodeCommit.If credentials are found, generate new credentials and store them in AWS KMS.C.   Configure Amazon Macie to scan for credentials in CodeCommit repositories. If credentials are found, trigger an AWS Lambda function to disable the credentials and notify the user.D.   Configure a CodeCommit trigger to invoke an AWS Lambda function to scan new code submissions for credentials. If credentials are found, disable them in AWS IAM and notify the user.Answer: CQUESTION 635A company is using AWS CodePipeline for the CI/CD of an application to an Amazon EC2 Auto Scaling group. All AWS resources are defined in AWS CloudFormation templates. The application artifacts are stored in an Amazon S3 bucket and deployed to the Auto Scaling group using instance user data scripts. As the application has become more complex, recent resource changes in the CloudFormation templates have caused unplanned downtime.How should a solutions architect improve the CI/CD pipeline to reduce the likelihood that changes in the templates will cause downtime?A.   Adapt the deployment scripts to detect and report CloudFormation error conditions when performing deployments. Write test plans for a testing team to execute in a non-production environment before approving the change for production.B.   Implement automated testing using AWS CodeBuild in a test environment. Use CloudFormation change sets to evaluate changes before deployment. Use AWS CodeDeploy to leverage blue/green deployment patterns to allow evaluations and the ability to revert changes, if needed.C.   Use plugins for the integrated development environment (IDE) to check the templates for errors, and use the AWS CLI to validate that the templates are correct. Adapt the deployment code to check for error

conditions and generate notifications on errors. Deploy to a test environment and execute a manual test plan before approving the change for production.D.   Use AWS CodeDeploy and a blue/green deployment pattern with CloudFormation to replace the user data deployment scripts. Have the operators log in to running instances and go through a manual test plan to verify the application is running as expected.Answer: DQUESTION 636A financial services company is moving to AWS and wants to enable developers to experiment and innovate while preventing access to production applications. The company has the following requirements:- Production workloads cannot be directly connected to the internet.- All workloads must be restricted to the us-west-2 and eu-central-1 Regions.- Notification should be sent when developer sandboxes exceed $500 in AWS spending monthly.Which combination of actions needs to be taken to create a multi-account structure that meets the company's requirements? (Choose three.)A.   Create accounts for each production workload within an organization in AWS Organizations. Place the production accounts within an organizational unit (OU). For each account, delete the default VPC.Create an SCP with a Deny rule for the attach an internet gateway and create a default VPC actions.Attach the SCP to the OU for the production accounts.B.   Create accounts for each production workload within an organization in AWS Organizations. Place the production accounts within an organizational unit (OU). Create an SCP with a Deny rule on the attach an internet gateway action. and create a default VPC actions. Create an SCP with a Deny rule to prevent use of the default VPC. Attach the SCPs to the OU for the production accounts.C.   Create a SCP containing a Deny Effect for cloudfront:*, iam:*, route53:*, and support:* with a StringNotEquals condition on an aws:RequestedRegion condition key with us-west-2 and eu-central-1 values. Attach the SCP to the organization's root.D.   Create an IAM permission boundary containing a Deny Effect for cloudfront:*, iam:*, route53:*, and support:* with a StringNotEquals condition on an aws:RequestedRegion condition key with us-west-2and eu-central-1 values. Attach the permission boundary to an IAM group containing the development and production users.E.   Create accounts for each development workload within an organization in AWS Organizations. Place the development accounts within an organizational unit (OU). Create a custom AWS Config rule to deactivate all IAM users when an account's monthly bill exceeds $500.F.   Create accounts for each development workload within an organization in AWS Organizations. Place the development accounts within an organizational unit (OU). Create a budget within AWS Budgets for each development account to monitor and report on monthly spending exceeding $500.Answer: ABCQUESTION 637A company is hosting a three-tier web application in an on-premises environment. Due to a recent surge in traffic that resulted in downtime and a significant financial impact, company management has ordered that the application be moved to AWS. The application is written in .NET and has a dependency on a MySQL database. A solutions architect must design a scalable and highly available solution to meet the demand of 200,000 daily users.Which steps should the solutions architect take to design an appropriate solution?A.   Use AWS Elastic Beanstalk to create a new application with a web server environment and an Amazon RDS MySQL Multi-AZ DB instance. The environment should launch a Network Load Balancer (NLB) in front of an Amazon EC2 Auto Scaling group in multiple Availability Zones. Use an Amazon Route 53 alias record to route traffic from the company's domain to the NLB.B.   Use AWS CloudFormation to launch a stack containing an Application Load Balancer (ALB) in front of an Amazon EC2 Auto Scaling group spanning three Availability Zones. The stack should launch a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a Retain deletion policy. Use an Amazon Route 53 alias record to route traffic from the company's domain to the ALB.C.   Use AWS Elastic Beanstalk to create an automatically scaling web server environment that spans two separate Regions with an Application Load Balancer (ALB) in each Region. Create a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a cross-Region read replica. Use Amazon Route 53 with a geoproximity routing policy to route traffic between the two Regions.D.   Use AWS CloudFormation to launch a stack containing an Application Load Balancer (ALB) in front of an Amazon ECS cluster of Spot instances spanning three Availability Zones. The stack should launch an Amazon RDS MySQL DB instance with a Snapshot deletion policy. Use an Amazon Route 53 alias record to route traffic from the company's domain to the ALB.Answer: AQUESTION 638A solutions architect is designing a publicly accessible web application that is on an Amazon CloudFront distribution with an Amazon S3 website endpoint as the origin. When the solution is deployed, the website returns an Error 403: Access Denied message.Which steps should the solutions architect take to correct the issue? (Choose two.)A.   Remove the S3 block public access option from the S3 bucket.B.   Remove the requester pays option from the S3 bucket. C.   Remove the origin access identity (OAI) from the CloudFront distribution.D.   Change the storage class from S3 Standard to S3 One Zone-Infrequent Access (S3 One Zone-IA).E.   Disable S3 object versioning.Answer: ACQUESTION 639A web application is hosted in a dedicated VPC that is connected to a company's on-premises data center over a Site-to-Site VPN connection. The application is accessible from the company network only. This is a temporary non-production application that is used during business hours. The workload is generally low with occasional surges.The application has an Amazon Aurora MySQL provisioned database cluster on the backend. The VPC has an internet gateway and a NAT gateways attached. The web servers are in private subnets in an Auto Scaling group behind an Elastic Load Balancer. The web servers also upload data to an Amazon S3 bucket

through the internet.A solutions architect needs to reduce operational costs and simplify the architecture.Which strategy should the solutions architect use?A.   Review the Auto Scaling group settings and ensure the scheduled actions are specified to operate the Amazon EC2 instances during business hours only. Use 3-year scheduled Reserved Instances for the web server EC2 instances. Detach the internet gateway and remove the NAT gateways from the VPC.Use an Aurora Serverless database and set up a VPC endpoint for the S3 bucket.B.   Review the Auto Scaling group settings and ensure the scheduled actions are specified to operate the Amazon EC2 instances during business hours only. Detach the internet gateway and remove the NAT gateways from the VPC. Use an Aurora Serverless database and set up a VPC endpoint for the S3 bucket, then update the network routing and security rules and policies related to the changes.C.   Review the Auto Scaling group settings and ensure the scheduled actions are specified to operate the Amazon EC2 instances during business hours only. Detach the internet gateway from the VPC, and use an Aurora Serverless database. Set up a VPC endpoint for the S3 bucket, then update the network routing and security rules and policies related to the changes.D.   Use 3-year scheduled Reserved Instances for the web server Amazon EC2 instances. Remove the NAT gateways from the VPC, and set up a VPC endpoint for the S3 bucket. Use Amazon CloudWatch and AWS Lambda to stop and start the Aurora DB cluster so it operates during business hours only.Update the network routing and security rules and policies related to the changes.Answer: CQUESTION 640A company plans to refactor a monolithic application into a modern application design deployed on AWS. The CI/CD pipeline needs to be upgraded to support the modern design for the application with the following requirements:- It should allow changes to be released several times every hour.- It should be able to roll back the changes as quickly as possible.Which design will meet these requirements?A.   Deploy a CI/CD pipeline that incorporates AMIs to contain the application and their configurations.Deploy the application by replacing Amazon EC2 instances.B.   Specify AWS Elastic Beanstalk to stage in a secondary environment as the deployment target for the CI/CD pipeline of the application. To deploy, swap the staging and production environment URLs.C.   Use AWS Systems Manager to re-provision the infrastructure for each deployment. Update the Amazon EC2 user data to pull the latest code artifact from Amazon S3 and use Amazon Route 53 weighted routing to point to the new environment.D.   Roll out the application updates as part of an Auto Scaling event using prebuilt AMIs. Use new versions of the AMIs to add instances, and phase out all instances that use the previous AMI version with the configured termination policy during a deployment event.Answer: AQUESTION 641A company currently has data hosted in an IBM Db2 database. A web application calls an API that runs stored procedures on the database to retrieve user information data that is read-only. This data is historical in nature and changes on a daily basis. When a user logs in to the application, this data needs to be retrieved within 3 seconds. Each time a user logs in, the stored procedures run. Users log in several times a day to check stock prices.Running this database has become cost-prohibitive due to Db2 CPU licensing. Performance goals are not being met. Timeouts from Db2 are common due to long-running queries.Which approach should a solutions architect take to migrate this solution to AWS?A.   Rehost the Db2 database in Amazon Fargate. Migrate all the data. Enable caching in Fargate. Refactor the API to use the Fargate Db2 database. Implement Amazon API Gateway and enable API caching.B.   Use AWS DMS to migrate data to Amazon DynamoDB using a continuous replication task. Refactor the API to use the DynamoDB data. Implement the refactored API in Amazon API Gateway and enable API caching.C.   Create a local cache on the mainframe to store query outputs. Use SFTP to sync to Amazon S3 on a daily basis. Refactor the API to use Amazon EFS.Implement Amazon API Gateway and enable API caching.D.   Extract data daily and copy the data to AWS Snowball for storage on Amazon S3. Sync daily. Refactor the API to use the S3 data. Implement Amazon API Gateway and enable API caching.Answer: AQUESTION 642A company is planning to deploy a new business analytics application that requires 10,000 hours of compute time each month. The compute resources can have flexible availability, but must be as cost- effective as possible. The company will also provide a reporting service to distribute analytics reports, which needs to run at all times.How should the solutions architect design a solution that meets these requirements?A. Deploy the reporting service on a Spot Fleet. Deploy the analytics application as a container in Amazon ECS with AWS Fargate as the compute option. Set the analytics application to use a custom metric with Service Auto Scaling.B.   Deploy the reporting service on an On-Demand Instance. Deploy the analytics application as a container in AWS Batch with AWS Fargate as the compute option. Set the analytics application to use a custom metric with Service Auto Scaling.C.   Deploy the reporting service as a container in Amazon ECS with AWS Fargate as the compute option.Deploy the analytics application on a Spot Fleet. Set the analytics application to use a custom metric with Amazon EC2 Auto Scaling applied to the Spot Fleet.D.   Deploy the reporting service as a container in Amazon ECS with AWS Fargate as the compute option.Deploy the analytics application on an On-Demand instance and purchase a Reserved Instance with a 3-year term. Set the analytics application to use a custom metric with Amazon EC2 Auto Scaling applied to the On-Demand instance.Answer: CQUESTION 643A 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. Launch instances from the AMIs created by AWS SMS. After initial testing, perform a final replication and create new instances from the updated AMIs.B.   Create an AWS CLI VM Import/Export script to migrate each virtual machine. Schedule the script to run incrementally to maintain changes in the application. Launch instances from the AMIs created by VM Import/Export. Once testing is done, rerun the script to do a final import and launch the instances from the AMIs.C.   Use AWS Server Migration Service (SMS) to upload the operating system volumes. Use the AWS CLI import-snapshot command for the data volumes. Launch instances from the AMIs created by AWS SMS and attach the data volumes to the instances. After initial testing, perform a final replication,launch new instances from the replicated AMIs, and attach the data volumes to the instances.D.   Use AWS Application Discovery Service and AWS Migration Hub to group the virtual machines as an application. Use the AWS CLI VM Import/Export script to import the virtual machines as AMIs.Schedule the script to run incrementally to maintain changes in the application. Launch instances from the AMIs. After initial testing, perform a final virtual machine import and launch new instances from the AMIs.Answer: BQUESTION 644An enterprise company's data science team wants to provide a safe, cost-effective way to provide easy access to Amazon SageMaker. The data scientists have limited AWS knowledge and need to be able to launch a Jupyter notebook instance. The notebook instance needs to have a preconfigured AWS KMS key to encrypt data at rest on the machine learning storage volume without exposing the complex setup requirements.Which approach will allow the company to set up a self-service mechanism for the data scientists to launch Jupyter notebooks in its AWS accounts with the LEAST amount of operational overhead?A.   Create a serverless front end using a static Amazon S3 website to allow the data scientists to request a Jupyter notebook instance by filling out a form. Use Amazon API Gateway to receive requests from the S3 website and trigger a central AWS Lambda function to make an API call to Amazon SageMaker that will launch a notebook instance with a preconfigured KMS key for the data scientists. Then call back to the front-end website to display the URL to the notebook instance.B.   Create an AWS CloudFormation template to launch a Jupyter notebook instance using the AWS::SageMaker::NotebookInstance resource type with a preconfigured KMS key. Add a user-friendly name to the CloudFormation template. Display the URL to the notebook using the Outputs section.Distribute the CloudFormation template to the data scientists using a shared Amazon S3 bucket.C.   Create an AWS CloudFormation template to launch a Jupyter notebook instance using the AWS::SageMaker::NotebookInstance resource type with a preconfigured KMS key. Simplify the parameter names, such as the instance size, by mapping them to Small, Large, and X-Large using the Mappings section in CloudFormation. Display the URL to the notebook using the Outputs section, then upload the template into an AWS Service Catalog product in the data scientist's portfolio, and share it with the data scientist IAM role.D.   Create an AWS CLI script that the data scientists can run locally. Provide step-by-step instructions about the parameters to be provided while executing the AWS CLI script to launch a Jupyter notebook with a preconfigured KMS key. Distribute the CLI script to the data scientists using a shared Amazon S3 bucket.Answer: BQUESTION 645A company is migrating its applications to AWS. The applications will be deployed to AWS accounts owned by business units. The company has several teams of developers who are responsible for the development and maintenance of all applications. The company is expecting rapid growth in the number of users.The company's chief technology officer has the following requirements:- Developers must launch the AWS infrastructure using AWS CloudFormation.- Developers must not be able to create resources outside of CloudFormation.- The solution must be able to scale to hundreds of AWS accounts. Which of the following would meet these requirements? (Choose two.)A.   Using CloudFormation, create an IAM role that can be assumed by CloudFormation that has permissions to create all the resources the company needs. Use CloudFormation StackSets to deploy this template to each AWS account.B.   In a central account, create an IAM role that can be assumed by developers, and attach a policy that allows interaction with CloudFormation. Modify the AssumeRolePolicyDocument action to allow the IAM role to be passed to CloudFormation.C.   Using CloudFormation, create an IAM role that can be assumed by developers, and attach policies that allow interaction with and passing a role to CloudFormation. Attach an inline policy to deny access to all other AWS services. Use CloudFormation StackSets to deploy this template to each AWS account.D.   Using CloudFormation, create an IAM role for each developer, and attach policies that allow interaction with CloudFormation. Use CloudFormation StackSets to deploy this template to each AWS account.E.   In a central AWS account, create an IAM role that can be assumed by CloudFormation that has permissions to create the resources the company requires. Create a CloudFormation stack policy that allows the IAM role to manage resources. Use CloudFormation StackSets to deploy the CloudFormation stack policy to each AWS account.Answer: ABQUESTION 646A media company has a static web application that is generated programmatically. The company has a build pipeline that generates HTML content that is uploaded to an Amazon S3 bucket served by Amazon CloudFront. The build pipeline runs inside a Build Account. The S3 bucket and CloudFront distribution are in a Distribution Account. The build pipeline uploads

the files to Amazon S3 using an IAM role in the Build Account. The S3 bucket has a bucket policy that only allows CloudFront to read objects using an origin access identity (OAI). During testing all attempts to access the application using the CloudFront URL result in an HTTP 403 Access Denied response.What should a solutions architect suggest to the company to allow access the objects in Amazon S3 through CloudFront?A.   Modify the S3 upload process in the Build Account to add the bucket-owner-full-control ACL to the objects at upload.B.   Create a new cross-account IAM role in the Distribution Account with write access to the S3 bucket.Modify the build pipeline to assume this role to upload the files to the Distribution Account.C.   Modify the S3 upload process in the Build Account to set the object owner to the Distribution Account.D.   Create a new IAM role in the Distribution Account with read access to the S3 bucket. Configure CloudFront to use this new role as its OAI. Modify the build pipeline to assume this role when uploading files from the Build Account.Answer: DQUESTION 647A 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? (Choose 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 clusters is launched across multiple Availability ZonesE.   Replace Amazon EFS win multiple Amazon EBS volumes in a RAID array.F. Replace Amazon EFS with Amazon FSx for Lustre.Answer: DEF