

## [2020-November-NewSAA-C02 VCE Exam Dumps SAA-C02 445Q Free Shared by Braindump2go[Q402-Q425]

2020/November Latest Braindump2go SAA-C02 Exam Dumps with PDF and VCE Free Updated Today! Following are some new

SAA-C02 Real Exam Questions!QUESTION 402A solutions architect is creating a new VPC design. There are two public subnet for the load balancer, two private subnets for web servers, and two private subnets for MySQL. The web servers use only HTTPS. The solutions architect has already created a security group for the load balancer allowing port 443 from 0.0.0.0/0. Company policy requires that each resource has the least access required to still be able to perform its tasks. Which additional configuration strategy should the solution architect use to meet these requirements?A. Create a security group for the web servers and allow port 443 from 0.0.0.0/0. Create a security group for the MySQL servers and allow port 3306 from the web servers security group.B. Create a network ACL for the web servers and allow port 443 from 0.0.0.0/0. Create a network ACL for the MySQL servers and allow port 3306 from the web servers security group.C. Create a security group for the web servers and allow port 443 from the load balancer. Create a security group for the MySQL servers and allow port 3306 from the web servers security group.D. Create a network ACL for the web servers and allow port 443 from the load balancer. Create a network ACL for the MySQL servers and allow port 3306 from the web servers security group. Answer: CQUESTION 403A company runs an application on an Amazon EC2 instance backed by Amazon Elastic Block Store (Amazon EBS). The instance needs to be available for 12 hours daily. The company wants to save costs by making the instance unavailable outside the window required for the application. However the contents of the instance's memory must be preserved whenever the instance is unavailable. What should a solutions architect do to meet this requirement?A. Stop the instance outside the application's availability window. Start up the instance again when required.B. Hibernate the instance outside the application's availability window. Start up the instance again when required.C. Use Auto Scaling to scale down the instance outside the application's availability window. Scale up the instance when required.D. Terminate the instance outside the application's availability window. Launch the instance by using a preconfigured Amazon Machine Image (AMI) when required. Answer: BQUESTION 404A company is migrating to the AWS Cloud. A file server is the first workload to migrate. Users must be able to access the file share using the Server Message Block (SMB) protocol. Which AWS managed service meets these requirements?A. Amazon EBSB. Amazon EC2C. Amazon FSxD. Amazon S3 Answer: BQUESTION 405A solutions architect needs to design a resilient solution for Windows users' home directories. The solution must provide fault tolerance, file-level backup and recovery, and access control, based upon the company's Active Directory. Which storage solution meets these requirements?A. Configure Amazon S3 to store the users' home directories. Join Amazon S3 to Active Directory.B. Configure a Multi-AZ file system with Amazon FSx for Windows File Server. Join Amazon FSx to Active Directory.C. Configure Amazon Elastic File System (Amazon EFS) for the users' home directories. Configure AWS Single Sign-On with Active Directory.D. Configure Amazon Elastic Block Store (Amazon EBS) to store the users' home directories. Configure AWS Single Sign-On with Active Directory. Answer: AQUESTION 406A company has a legacy application that processes data in two parts. The second part of the process takes longer than the first, so the company has decided to rewrite the application as two microservices running on Amazon ECS that can scale independently. How should a solutions architect integrate the microservices?A. Implement code in microservice 1 to send data to an Amazon S3 bucket. Use S3 event notifications to invoke microservice 2.B. Implement code in microservice 1 to publish data to an Amazon SNS topic. Implement code in microservice 2 to subscribe to this topic.C. Implement code in microservice 1 to send data to Amazon Kinesis Data Firehose. Implement code in microservice 2 to read from Kinesis Data Firehose.D. Implement code in microservice 1 to send data to an Amazon SQS queue. Implement code in microservice 2 to process messages from the queue. Answer: AQUESTION 407A company hosts its application using Amazon Elastic Container Service (Amazon ECS) and wants to ensure high availability. The company wants to be able to deploy updates to its application even if nodes in one Availability Zone are not accessible. The expected request volume for the application is 100 requests per second, and each container task is able to serve at least 60 requests per second. The company set up Amazon ECS with a rolling update deployment type with the minimum healthy percent parameter set to 50% and the maximum percent set to 100%. Which configuration of tasks and Availability Zones meets these requirements?A. Deploy the application across two Availability Zones, with one task in each Availability Zone.B. Deploy the application across two Availability Zones, with two tasks in each Availability Zone.C. Deploy the application across three Availability Zones, with one task in each Availability Zone.D. Deploy the application across three Availability Zones, with two tasks in each Availability Zone. Answer: AQUESTION 408A web application runs on Amazon EC2 instances behind an Application Load Balancer. The application allows users to create custom reports of historical weather data. Generating a report can take up to 5 minutes. These long-running requests use many of the available incoming connections, making the system unresponsive to other users. How can a solutions architect make the system more

responsive?A. Use Amazon SOS with AWS Lambda to generate reports.B. Increase the Idle timeout on the Application Load Balancer to 5 minutes.C. Update the client-side application code to increase its request timeout to 5 minutes.D. Publish the reports to Amazon S3 and use Amazon CloudFront for downloading to the user.  
Answer: A  
QUESTION 409A company is planning to use Amazon S3 to store images uploaded by its users. The images must be encrypted at rest in Amazon S3. The company does not want to spend time managing and rotating the keys, but it does want to control who can access those keys. What should a solutions architect use to accomplish this?  
A. Server-Side Encryption with keys stored in an S3 bucket  
B. Server-Side Encryption with Customer-Provided Keys (SSE-C)  
C. Server-Side Encryption with Amazon S3-Managed Keys (SSE-S3)  
D. Server-Side Encryption with AWS KMS-Managed Keys (SSE-KMS)  
Answer: D  
QUESTION 410A company's application is running on Amazon EC2 instances within an Auto Scaling group behind an Elastic Load Balancer. Based on the application's history, the company anticipates a spike in traffic during a holiday each year. A solutions architect must design a strategy to ensure that the Auto Scaling group proactively increases capacity to minimize any performance impact on application users. Which solution will meet these requirements?  
A. Create an Amazon CloudWatch alarm to scale up the EC2 instances when CPU utilization exceeds 90%.  
B. Create a recurring scheduled action to scale up the Auto Scaling group before the expected period of peak demand.  
C. Increase the minimum and maximum number of EC2 instances in the Auto Scaling group during the peak demand period.  
D. Configure an Amazon Simple Notification Service (Amazon SNS) notification to send alerts when there are autoscaling:EC2\_INSTANCE\_LAUNCH events.  
Answer: B  
QUESTION 411A website runs a web application that receives a burst of traffic each day at noon. The users upload new pictures and content daily, but have been complaining of timeouts. The architecture uses Amazon EC2 Auto Scaling groups, and the custom application consistently takes 1 minute to initiate upon boot up before responding to user requests. How should a solutions architect redesign the architecture to better respond to changing traffic?  
A. Configure a Network Load Balancer with a slow start configuration.  
B. Configure AWS ElastiCache for Redis to offload direct requests to the servers.  
C. Configure an Auto Scaling step scaling policy with an instance warmup condition.  
D. Configure Amazon CloudFront to use an Application Load Balancer as the origin.  
Answer: B  
QUESTION 412A solutions architect needs to design a managed storage solution for a company's application that includes high-performance machine learning. This application runs on AWS Fargate, and the connected storage needs to have concurrent access to files and deliver high performance. Which storage option should the solutions architect recommend?  
A. Create an Amazon S3 bucket for the application and establish an IAM role for Fargate to communicate with Amazon S3.  
B. Create an Amazon FSx for Lustre file share and establish an IAM role that allows Fargate to communicate with FSx for Lustre.  
C. Create an Amazon Elastic File System (Amazon EFS) file share and establish an IAM role that allows Fargate to communicate with Amazon EFS.  
D. Create an Amazon Elastic Block Store (Amazon EBS) volume for the application and establish an IAM role that allows Fargate to communicate with Amazon EBS.  
Answer: B  
QUESTION 413A company is launching an e-commerce website on AWS. This website is built with a three-tier architecture that includes a MySQL database. In a Multi-AZ deployment of Amazon Aurora MySQL. The website application must be highly available and will initially be launched in an AWS Region with three Availability Zones. The application produces a metric that describes the load the application experiences. Which solution meets these requirements?  
A. Configure an Application Load Balancer (ALB) with Amazon EC2 Auto Scaling behind the ALB with scheduled scaling.  
B. Configure an Application Load Balancer (ALB) and Amazon EC2 Auto Scaling behind the ALB with a simple scaling policy.  
C. Configure a Network Load Balancer (NLB) and launch a Spot Fleet with Amazon EC2 Auto Scaling behind the NLB.  
D. Configure an Application Load Balancer (ALB) and Amazon EC2 Auto Scaling behind the ALB with a target tracking scaling policy.  
Answer: B  
QUESTION 414A company is designing an internet-facing web application. The application runs on Amazon EC2 for Linux-based instances that store sensitive user data in Amazon RDS MySQL Multi-AZ DB instances. The EC2 instances are in public subnets, and the RDS DB instances are in private subnets. The security team has mandated that the DB instances be secured against web-based attacks. What should a solutions architect recommend?  
A. Ensure the EC2 instances are part of an Auto Scaling group and are behind an Application Load Balancer. Configure the EC2 instance iptables rules to drop suspicious web traffic. Create a security group for the DB instances. Configure the RDS security group to only allow port 3306 inbound from the individual EC2 instances.  
B. Ensure the EC2 instances are part of an Auto Scaling group and are behind an Application Load Balancer. Move DB instances to the same subnets that EC2 instances are located in. Create a security group for the DB instances. Configure the RDS security group to only allow port 3306 inbound from the individual EC2 instances.  
C. Ensure the EC2 instances are part of an Auto Scaling group and are behind an Application Load Balancer. Use AWS WAF to monitor inbound web traffic for threats. Create a security group for the web application servers and a security group for the DB instances. Configure the RDS security group to only allow port 3306 inbound from the web application server security group.  
D. Ensure the EC2 instances are part of an Auto Scaling group and are behind an Application Load Balancer. Use AWS WAF to monitor inbound web traffic for threats. Configure the Auto Scaling group to

automatically create new DB instances under heavy traffic. Create a security group for the RDS DB instances. Configure the RDS security group to only allow port 3306 inbound. Answer: D

QUESTION 415A start-up company has a web application based in the us-east-1 Region with multiple Amazon EC2 instances running behind an Application Load Balancer across multiple Availability Zones. As the company's user base grows in the us-west-1 Region, it needs 3 solution with low latency and high availability. What should a solutions architect do to accomplish this? A. Provision EC2 instances in us-west-1. Switch the Application Load Balancer to a Network Load Balancer to achieve cross-Region load balancing. B. Provision EC2 instances and an Application Load Balancer in us-west-1. Make the load balancer distribute the traffic based on the location of the request. C. Provision EC2 instances and configure an Application Load Balancer in us-west-1. Create an accelerator in AWS Global Accelerator that uses an endpoint group that includes the load balancer endpoints in both Regions. D. Provision EC2 instances and configure an Application Load Balancer in us-west-1. Configure Amazon Route 53 with a weighted routing policy. Create alias records in Route 53 that point to the Application Load Balancer. Answer: C

Explanation: <https://aws.amazon.com/global-accelerator/faqs/>

QUESTION 416A company has a custom application running on an Amazon EC2 instance that: - Reads a large amount of data from Amazon S3 - Performs a multi-stage analysis. - Writes the results to Amazon DynamoDB. The application writes a significant number of large, temporary files during the multi-stage analysis. The process performance depends on the temporary storage performance. What would be the fastest storage option for holding the temporary files? A. Multiple Amazon S3 buckets with Transfer Acceleration for storage. B. Multiple Amazon EBS drives with Provisioned IOPS and EBS optimization. C. Multiple Amazon EFS volumes using the Network File System version 4.1 (NFSv4.1) protocol. D. Multiple instance store volumes with software RAID 0. Answer: D

QUESTION 417A company built a food ordering application that captures user data and stores it for future analysis. The application's static front end is deployed on an Amazon EC2 instance. The front-end application sends the requests to the backend application running on separate EC2 instance. The backend application then stores the data in Amazon RDS. What should a solutions architect do to decouple the architecture and make it scalable? A. Use Amazon S3 to serve the front-end application, which sends requests to Amazon EC2 to execute the backend application. The backend application will process and store the data in Amazon RDS. B. Use Amazon S3 to serve the front-end application and write requests to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe Amazon EC2 instances to the HTTP/HTTPS endpoint of the topic, and process and store the data in Amazon RDS. C. Use an EC2 instance to serve the front end and write requests to an Amazon SQS queue. Place the backend instance in an Auto Scaling group, and scale based on the queue depth to process and store the data in Amazon RDS. D. Use Amazon S3 to serve the static front-end application and send requests to Amazon API Gateway which writes the requests to an Amazon SQS queue. Place the backend instances in an Auto Scaling group, and scale based on the queue depth to process and store the data in Amazon RDS. Answer: D

QUESTION 418A company has an on-premises application that collects data and stores it to an on-premises NFS server. The company recently set up a 10 Gbps AWS Direct Connect connection. The company is running out of storage capacity on premises. The company needs to migrate the application data from on premises to the AWS Cloud while maintaining low-latency access to the data from the on-premises application. What should a solutions architect do to meet these requirements? A. Deploy AWS Storage Gateway for the application data, and use the file gateway to store the data in Amazon S3. Connect the on-premises application servers to the file gateway using NFS. B. Attach an Amazon Elastic File System (Amazon EFS) file system to the NFS server, and copy the application data to the EFS file system. Then connect the on-premises application to Amazon EFS. C. Configure AWS Storage Gateway as a volume gateway. Make the application data available to the on-premises application from the NFS server and with Amazon Elastic Block Store (Amazon EBS) snapshots. D. Create an AWS DataSync agent with the NFS server as the source location and an Amazon Elastic File System (Amazon EFS) file system as the destination for application data transfer. Connect the on-premises application to the EFS file system. Answer: A

QUESTION 419A company wants to migrate a high performance computing (HPC) application and data from on-premises to the AWS Cloud. The company uses tiered storage on-premises with high-performance parallel storage to support the application during periodic runs of the application, and more economical cold storage to hold the data when the application is not actively running. Which combination of solutions should a solutions architect recommend to support the storage needs of the application? (Select TWO) A. Amazon S3 for cold data storage. B. Amazon EFS for cold data storage. C. Amazon S3 for high-performance parallel storage. D. Amazon FSx for Lustre for high-performance parallel storage. E. Amazon FSx for Windows for high-performance parallel storage. Answer: ADE

Explanation: <https://aws.amazon.com/fsx/lustre/>

Amazon FSx for Lustre makes it easy and cost effective to launch and run the world's most popular high-performance file system. Use it for workloads where speed matters, such as machine learning, high performance computing (HPC), video processing, and financial modeling.

QUESTION 420A software vendor is deploying a new software-as-a-service (SaaS) solution that will be utilized by many AWS users. The service is hosted in a VPC behind a Network Load Balancer. The software vendor wants to provide access to this service to users with the least amount of administrative

overhead and without exposing the service to the public internet. What should a solutions architect do to accomplish this goal? A. Create a peering VPC connection from each user's VPC to the software vendor's VPC. B. Deploy a transit VPC in the software vendor's AWS account. Create a VPN connection with each user account. C. Connect the service in the VPC with an AWS PrivateLink endpoint. Have users subscribe to the endpoint. D. Deploy a transit VPC in the software vendor's AWS account. Create an AWS Direct Connect connection with each user account. Answer: C

QUESTION 421 A company uses Amazon S3 to store its confidential audit documents. The S3 bucket uses bucket policies to restrict access to audit team IAM user credentials according to the principle of least privilege. Company managers are worried about accidental deletion of documents in the S3 bucket and want a more secure solution. What should a solutions architect do to secure the audit documents? A. Enable the versioning and MFA Delete features on the S3 bucket. B. Enable multi-factor authentication (MFA) on the IAM user credentials for each audit team IAM user account. C. Add an S3 Lifecycle policy to the audit team's IAM user accounts to deny the s3:DeleteObject action during audit dates. D. Use AWS Key Management Service (AWS KMS) to encrypt the S3 bucket and restrict audit team IAM user accounts from accessing the KMS key. Answer: A

QUESTION 422 A solutions architect is helping a developer design a new ecommerce shopping cart application using AWS services. The developer is unsure of the current database schema and expects to make changes as the ecommerce site grows. The solution needs to be highly resilient and capable of automatically scaling read and write capacity. Which database solution meets these requirements? A. Amazon Aurora PostgreSQL. B. Amazon DynamoDB with on-demand enabled. C. Amazon DynamoDB with DynamoDB Streams enabled. D. Amazon SQS and Amazon Aurora PostgreSQL. Answer: B

QUESTION 423 A company is seeing access requests by some suspicious IP addresses. The security team discovers the requests are from different IP addresses under the same CIDR range. What should a solutions architect recommend to the team? A. Add a rule in the inbound table of the security group to deny the traffic from that CIDR range. B. Add a rule in the outbound table of the security group to deny the traffic from that CIDR range. C. Add a deny rule in the inbound table of the network ACL with a lower rule number than other rules. D. Add a deny rule in the outbound table of the network ACL with a lower rule number than other rules. Answer: C

QUESTION 424 A company wants to run a hybrid workload for data processing. The data needs to be accessed by on-premises applications for local data processing using an NFS protocol, and must also be accessible from the AWS Cloud for further analytics and batch processing. Which solution will meet these requirements? A. Use an AWS Storage Gateway file gateway to provide file storage to AWS. Then perform analytics on the data in the AWS Cloud. B. Use an AWS Storage Gateway tape gateway to copy the backup of the local data to AWS. Then perform analytics on this data in the AWS Cloud. C. Use an AWS Storage Gateway volume gateway in a stored volume configuration to regularly take snapshots of the local data, then copy the data to AWS. D. Use an AWS Storage Gateway volume gateway in a cached volume configuration to back up all the local storage in the AWS Cloud, then perform analytics on this data in the cloud. Answer: C

Explanation:  
<https://docs.aws.amazon.com/storagegateway/latest/userguide/WhatIsStorageGateway.html>

QUESTION 425 A solutions architect is designing a solution that requires frequent updates to a website that is hosted on Amazon S3 with versioning enabled. For compliance reasons, older versions of the objects will not be accessed frequently and will need to be deleted after 2 years. What should the solutions architect recommend to meet these requirements at the LOWEST cost? A. Use S3 batch operations to replace object tags. Expire the objects based on the modified tags. B. Configure an S3 Lifecycle policy to transition older versions of objects to S3 Glacier. Expire the objects after 2 years. C. Enable S3 Event Notifications on the bucket that sends older objects to the Amazon Simple Queue Service (Amazon SQS) queue for further processing. D. Replicate older object versions to a new bucket. Use an S3 Lifecycle policy to expire the objects in the new bucket after 2 years. Answer: B

Resources From: 1. 2020 Latest Braindump2go SAA-C02 Exam Dumps (PDF & VCE) Free Share: <https://www.braindump2go.com/saa-c02.html> 2. 2020 Latest Braindump2go SAA-C02 PDF and SAA-C02 VCE Dumps Free Share: [https://drive.google.com/drive/folders/1\\_5IK3H\\_eM74C6AKwU7sKaLn1rm8xTfm?usp=sharing](https://drive.google.com/drive/folders/1_5IK3H_eM74C6AKwU7sKaLn1rm8xTfm?usp=sharing) 3. 2020 Free Braindump2go SAA-C02 PDF Download: [https://www.braindump2go.com/free-online-pdf/SAA-C02-Dumps\(416-430\).pdf](https://www.braindump2go.com/free-online-pdf/SAA-C02-Dumps(416-430).pdf) [https://www.braindump2go.com/free-online-pdf/SAA-C02-PDF-Dumps\(402-415\).pdf](https://www.braindump2go.com/free-online-pdf/SAA-C02-PDF-Dumps(402-415).pdf) [https://www.braindump2go.com/free-online-pdf/SAA-C02-VCE-Dumps\(431-445\).pdf](https://www.braindump2go.com/free-online-pdf/SAA-C02-VCE-Dumps(431-445).pdf) Free Resources from Braindump2go, We Devoted to Helping You 100% Pass All Exams!