

## [September-2022Real SAA-C03 VCE SAA-C03 210Q Free Download in Braindump2go[Q95-Q125

September/2022 Latest Braindump2go SAA-C03 Exam Dumps with PDF and VCE Free Updated Today! Following are some new Braindump2go SAA-C03 Real Exam Questions!QUESTION 95A company hosts its product information webpages on AWS. The existing solution uses multiple Amazon EC2 instances behind an Application Load Balancer in an Auto Scaling group. The website also uses a custom DNS name and communicates with HTTPS only using a dedicated SSL certificate. The company is planning a new product launch and wants to be sure that users from around the world have the best possible experience on the new website. What should a solutions architect do to meet these requirements?

A. Redesign the application to use Amazon CloudFront  
B. Redesign the application to use AWS Elastic Beanstalk  
C. Redesign the application to use a Network Load Balancer  
D. Redesign the application to use Amazon S3 static website hosting  
Answer: AExplanation: CloudFront can help provide the best experience for global users. CloudFront integrates seamlessly with ALB and provides an option to use custom DNS and SSL certs.  
QUESTION 96A company has 150 TB of archived image data stored on-premises that needs to be moved to the AWS Cloud within the next month. The company's current network connection allows up to 100 Mbps uploads for this purpose during the night only. What is the MOST cost-effective mechanism to move this data and meet the migration deadline?

A. Use AWS Snowmobile to ship the data to AWS.  
B. Order multiple AWS Snowball devices to ship the data to AWS.  
C. Enable Amazon S3 Transfer Acceleration and securely upload the data.  
D. Create an Amazon S3 VPC endpoint and establish a VPN to upload the data  
Answer: BExplanation: eg. 6 hrs night \* 6 hrs \* 60 min/hr = 360 min \* 60 sec/min = 21600 sec \* 100 Mbps \* 21600 s = 2160000 Mbit/s = 2160 Gb or 2.1 TB can only be done. So, for 150 TB, we can use 2 X Snowball Edge Storage Optimized devices. Size of Snowball Edge Storage Optimized device = 80 TB. Size of Snowball Edge Compute Optimized device = 40 TB. Size of Snowcone = 8 TB. Size of Snowmobile = 100 PB (1 PB = 1000 TB).  
Q: How should I choose between Snowmobile and Snowball? To migrate large datasets of 10PB or more in a single location, you should use Snowmobile. For datasets less than 10PB or distributed in multiple locations, you should use Snowball. In addition, you should evaluate the amount of available bandwidth in your network backbone. If you have a high speed backbone with hundreds of Gb/s of spare throughput, then you can use Snowmobile to migrate the large datasets all at once. If you have limited bandwidth on your backbone, you should consider using multiple Snowballs to migrate the data incrementally.

QUESTION 97A company hosts its web application on AWS using seven Amazon EC2 instances. The company requires that the IP addresses of all healthy EC2 instances be returned in response to DNS queries. Which policy should be used to meet this requirement?

A. Simple routing policy  
B. Latency routing policy  
C. Multivalue routing policy  
D. Geolocation routing policy  
Answer: CExplanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/multivalue-versus-simple-policies/> Use a multivalue answer routing policy to help distribute DNS responses across multiple resources. For example, use multivalue answer routing when you want to associate your routing records with a Route 53 health check.

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html#routing-policy-multivalue>QUESTION 98A company wants to use AWS Systems Manager to manage a fleet of Amazon EC2 instances. According to the company's security requirements, no EC2 instances can have internet access. A solutions architect needs to design network connectivity from the EC2 instances to Systems Manager while fulfilling this security obligation. Which solution will meet these requirements?

A. Deploy the EC2 instances into a private subnet with no route to the internet.  
B. Configure an interface VPC endpoint for Systems Manager. Update routes to use the endpoint.  
C. Deploy a NAT gateway into a public subnet. Configure private subnets with a default route to the NAT gateway.  
D. Deploy an internet gateway. Configure a network ACL to deny traffic to all destinations except Systems Manager.  
Answer: BExplanation: VPC Peering connectionsVPC interface endpoints can be accessed through both intra-Region and inter-Region VPC peering connections. VPC Gateway Endpoint connections can't be extended out of a VPC. Resources on the other side of a VPC peering connection in your VPC can't use the gateway endpoint to communicate with resources in the gateway endpoint service. Reference: <https://docs.aws.amazon.com/systems-manager/latest/userguide/setup-create-vpc.html>  
QUESTION 99A company needs to build a reporting solution on AWS. The solution must support SQL queries that data analysts run on the data. The data analysts will run fewer than 10 total queries each day. The company generates 3 GB of new data daily in an on-premises relational database. This data needs to be transferred to AWS to perform reporting tasks. What should a solutions architect recommend to meet these requirements at the LOWEST cost?

A. Use AWS Database Migration Service (AWS DMS) to replicate the data from the on-premises database into Amazon S3. Use Amazon Athena to query the data.  
B. Use an Amazon Kinesis Data Firehose delivery stream to deliver the data into an Amazon Elasticsearch Service (Amazon ES) cluster. Run the queries in Amazon ES.  
C. Export a daily copy of the data from the on-premises database. Use an AWS Storage Gateway file gateway to store and copy the export into Amazon S3. Use an Amazon EMR cluster to query the data.  
D. Use AWS Database Migration Service (AWS DMS)

to replicate the data from the on-premises database and load it into an Amazon Redshift cluster. Use the Amazon Redshift cluster to query the data. Answer: DExplanation: [https://docs.aws.amazon.com/dms/latest/userguide/CHAP\\_Target.Redshift.html](https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.Redshift.html) AWS DMS cannot migrate or replicate changes to a schema with a name that begins with underscore (\_). If you have schemas that have a name that begins with an underscore, use mapping transformations to rename the schema on the target. Amazon Redshift doesn't support VARCHARs larger than 64 KB. LOBs from traditional databases can't be stored in Amazon Redshift. Applying a DELETE statement to a table with a multi-column primary key is not supported when any of the primary key column names use a reserved word. Go here to see a list of Amazon Redshift reserved words. You may experience performance issues if your source system performs UPDATE operations on the primary key of a source table. These performance issues occur when applying changes to the target. This is because UPDATE (and DELETE) operations depend on the primary key value to identify the target row. If you update the primary key of a source table, your task log will contain messages like the following: Update on table 1 changes PK to a PK that was previously updated in the same bulk update. DMS doesn't support custom DNS names when configuring an endpoint for a Redshift cluster, and you need to use the Amazon provided DNS name. Since the Amazon Redshift cluster must be in the same AWS account and Region as the replication instance, validation fails if you use a custom DNS endpoint. QUESTION 100 A company wants to monitor its AWS costs for financial review. The cloud operations team is designing an architecture in the AWS Organizations management account to query AWS Cost and Usage Reports for all member accounts. The team must run this query once a month and provide a detailed analysis of the bill. Which solution is the MOST scalable and cost-effective way to meet these requirements? A. Enable Cost and Usage Reports in the management account. Deliver reports to Amazon Kinesis. Use Amazon EMR for analysis. B. Enable Cost and Usage Reports in the management account. Deliver the reports to Amazon S3. Use Amazon Athena for analysis. C. Enable Cost and Usage Reports for member accounts. Deliver the reports to Amazon S3. Use Amazon Redshift for analysis. D. Enable Cost and Usage Reports for member accounts. Deliver the reports to Amazon Kinesis. Use Amazon QuickSight for analysis. Answer: CExplanation: <https://docs.aws.amazon.com/cur/latest/userguide/what-is-cur.html> If you are an administrator of an AWS Organizations management account and do not want any of the member accounts in your Organization to set-up a CUR you can do one of the following: (Recommended) If you've opted into Organizations with all features enabled, you can apply a Service Control Policy (SCP). Note that SCPs only apply to member accounts and if you want to restrict any IAM users associated with the management account from setting up a CUR, you'll need to adjust their specific IAM permissions. SCPs also are not retroactive, so they will not de-activate any CURs a member account may have set-up prior to the SCP being applied. Submit a customer support case to block access to billing data in the Billing console for member accounts. This is a list of organizations where the payer account prevents member accounts in its organization from viewing billing data on the Bills and Invoices pages. This also prevents those accounts from setting up Cost and Usage Reports. This option is only available for organizations without all features enabled. Please note that if you have already opted into this to prevent member accounts from viewing bills and invoices in the Billing Console, you do not need to request this access again. Those same member accounts will also be prevented from setting up a Cost and Usage Report. QUESTION 101 A company collects temperature, humidity, and atmospheric pressure data in cities across multiple continents. The average volume of data collected per site each day is 500 GB. Each site has a high-speed internet connection. The company's weather forecasting applications are based in a single Region and analyze the data daily. What is the FASTEST way to aggregate data from all of these global sites? A. Enable Amazon S3 Transfer Acceleration on the destination bucket. Use multipart uploads to directly upload site data to the destination bucket. B. Upload site data to an Amazon S3 bucket in the closest AWS Region. Use S3 cross-Region replication to copy objects to the destination bucket. C. Schedule AWS Snowball jobs daily to transfer data to the closest AWS Region. Use S3 cross-Region replication to copy objects to the destination bucket. D. Upload the data to an Amazon EC2 instance in the closest Region. Store the data in an Amazon Elastic Block Store (Amazon EBS) volume. Once a day take an EBS snapshot and copy it to the centralized Region. Restore the EBS volume in the centralized Region and run an analysis on the data daily. Answer: AExplanation: You might want to use Transfer Acceleration on a bucket for various reasons, including the following: - You have customers that upload to a centralized bucket from all over the world. - You transfer gigabytes to terabytes of data on a regular basis across continents. - You are unable to utilize all of your available bandwidth over the Internet when uploading to Amazon S3. <https://docs.aws.amazon.com/AmazonS3/latest/dev/transfer-acceleration.html> [https://aws.amazon.com/s3/transfer-acceleration/#:~:text=S3%20Transfer%20Acceleration%20\(S3TA\)%20reduces,to%20S3%20for%20remote%20applications](https://aws.amazon.com/s3/transfer-acceleration/#:~:text=S3%20Transfer%20Acceleration%20(S3TA)%20reduces,to%20S3%20for%20remote%20applications) Amazon S3 Transfer Acceleration can speed up content transfers to and from Amazon S3 by as much as 50-500% for long-distance transfer of larger objects. Customers who have either web or mobile applications with widespread users or applications hosted far away from their S3 bucket can experience long and variable upload and download speeds over the Internet. <https://docs.aws.amazon.com/AmazonS3/latest/userguide/mpuoverview.html> Improved throughput - You can upload parts in parallel to improve throughput. QUESTION 102 A company needs the ability to analyze the log files of its proprietary

application. The logs are stored in JSON format in an Amazon S3 bucket. Queries will be simple and will run on-demand. A solutions architect needs to perform the analysis with minimal changes to the existing architecture. What should the solutions architect do to meet these requirements with the LEAST amount of operational overhead?

A. Use Amazon Redshift to load all the content into one place and run the SQL queries as needed.

B. Use Amazon CloudWatch Logs to store the logs and run SQL queries as needed from the Amazon CloudWatch console.

C. Use Amazon Athena directly with Amazon S3 to run the queries as needed.

D. Use AWS Glue to catalog the logs.

Use a transient Apache Spark cluster on Amazon EMR to run the SQL queries as needed.

Answer: C

Explanation: Amazon Athena can be used to query JSON in S3.

QUESTION 103

A company uses AWS Organizations to manage multiple AWS accounts for different departments. The management account has an Amazon S3 bucket that contains project reports. The company wants to limit access to this S3 bucket to only users of accounts within the organization in AWS Organizations. Which solution meets these requirements with the LEAST amount of operational overhead?

A. Add the aws:PrincipalOrgID global condition key with a reference to the organization ID to the S3 bucket policy.

B. Create an organizational unit (OU) for each department. Add the aws:PrincipalOrgPaths global condition key to the S3 bucket policy.

C. Use AWS CloudTrail to monitor the CreateAccount, InviteAccountToOrganization, LeaveOrganization, and RemoveAccountFromOrganization events. Update the S3 bucket policy accordingly.

D. Tag each user that needs access to the S3 bucket. Add the aws:PrincipalTag global condition key to the S3 bucket policy.

Answer: A

Explanation:

<https://aws.amazon.com/blogs/security/control-access-to-aws-resources-by-using-the-aws-organization-of-iam-principals/>

The aws:PrincipalOrgID global key provides an alternative to listing all the account IDs for all AWS accounts in an organization. For example, the following Amazon S3 bucket policy allows members of any account in the XXX organization to add an object into the examtopics bucket.

```
{"Version": "2020-09-10", "Statement": { "Sid": "AllowPutObject", "Effect": "Allow", "Principal": "*", "Action": "s3:PutObject", "Resource": "arn:aws:s3:::examtopics/*", "Condition": { "StringEquals": { "aws:PrincipalOrgID": "XXX" } } }}
```

[https://docs.aws.amazon.com/IAM/latest/UserGuide/reference\\_policies\\_condition-keys.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_condition-keys.html)

QUESTION 104

An application runs on an Amazon EC2 instance in a VPC. The application processes logs that are stored in an Amazon S3 bucket. The EC2 instance needs to access the S3 bucket without connectivity to the internet. Which solution will provide private network connectivity to Amazon S3?

A. Create a gateway VPC endpoint to the S3 bucket.

B. Stream the logs to Amazon CloudWatch Logs. Export the logs to the S3 bucket.

C. Create an instance profile on Amazon EC2 to allow S3 access.

D. Create an Amazon API Gateway API with a private link to access the S3 endpoint.

Answer: A

QUESTION 105

A company is hosting a web application on AWS using a single Amazon EC2 instance that stores user-uploaded documents in an Amazon EBS volume. For better scalability and availability, the company duplicated the architecture and created a second EC2 instance and EBS volume in another Availability Zone placing both behind an Application Load Balancer. After completing this change, users reported that, each time they refreshed the website, they could see one subset of their documents or the other, but never all of the documents at the same time. What should a solutions architect propose to ensure users see all of their documents at once?

A. Copy the data so both EBS volumes contain all the documents.

B. Configure the Application Load Balancer to direct a user to the server with the documents.

C. Copy the data from both EBS volumes to Amazon EFS. Modify the application to save new documents to Amazon EFS.

D. Configure the Application Load Balancer to send the request to both servers. Return each document from the correct server.

Answer: C

Explanation: Amazon EFS provides file storage in the AWS Cloud. With Amazon EFS, you can create a file system, mount the file system on an Amazon EC2 instance, and then read and write data to and from your file system. You can mount an Amazon EFS file system in your VPC, through the Network File System versions 4.0 and 4.1 (NFSv4) protocol. We recommend using a current generation Linux NFSv4.1 client, such as those found in the latest Amazon Linux, Redhat, and Ubuntu AMIs, in conjunction with the Amazon EFS Mount Helper. For instructions, see Using the amazon-efs-utils Tools. For a list of Amazon EC2 Linux Amazon Machine Images (AMIs) that support this protocol, see NFS Support. For some AMIs, you'll need to install an NFS client to mount your file system on your Amazon EC2 instance. For instructions, see Installing the NFS Client. You can access your Amazon EFS file system concurrently from multiple NFS clients, so applications that scale beyond a single connection can access a file system. Amazon EC2 instances running in multiple Availability Zones within the same AWS Region can access the file system, so that many users can access and share a common data source.

QUESTION 106

A company uses NFS to store large video files in on-premises network attached storage. Each video file ranges in size from 1MB to 500 GB. The total storage is 70 TB and is no longer growing. The company decides to migrate the video files to Amazon S3. The company must migrate the video files as soon as possible while using the least possible network bandwidth. Which solution will meet these requirements?

A. Create an S3 bucket. Create an IAM role that has permissions to write to the S3 bucket. Use the AWS CLI to copy all files locally to the S3 bucket.

B. Create an AWS Snowball Edge job. Receive a Snowball Edge device on premises. Use the Snowball Edge client to transfer data to the device. Return the device so that AWS can import the data into Amazon S3.

C. Deploy an S3 File Gateway on premises. Create a public service

endpoint to connect to the S3 File Gateway.Create an S3 bucket.Create a new NFS file share on the S3 File Gateway.Point the new file share to the S3 bucket. Transfer the data from the existing NFS file share to the S3 File Gateway.D. Set up an AWS Direct Connect connection between the on-premises network and AWS. Deploy an S3 File Gateway on premises. Create a public virtual interface (VIF) to connect to the S3 File Gateway.Create an S3 bucket. Create a new NFS file share on the S3 File Gateway. Point the new file share to the S3 bucket. Transfer the data from the existing NFS file share to the S3 File Gateway.Answer: CQUESTION 107A company has an application that ingests incoming messages. These messages are then quickly consumed by dozens of other applications and microservices. The number of messages varies drastically and sometimes spikes as high as 100,000 each second. The company wants to decouple the solution and increase scalability.Which solution meets these requirements?A. Persist the messages to Amazon Kinesis Data Analytics. All the applications will read and process the messages.B. Deploy the application on Amazon EC2 instances in an Auto Scaling group, which scales the number of EC2 instances based on CPU metrics.C. Write the messages to Amazon Kinesis Data Streams with a single shard. All applications will read from the stream and process the messages.D. Publish the messages to an Amazon Simple Notification Service (Amazon SNS) topic with one or more Amazon Simple Queue Service (Amazon SQS) subscriptions. All applications then process the messages from the queues.Answer: DExplanation:"SNS Standard Topic"Maximum throughput: Standard topics support a nearly unlimited number of messages per second.

<https://aws.amazon.com/sns/features/>"SQS Standard Queue"Unlimited Throughput: Standard queues support a nearly unlimited number of transactions per second (TPS) per API action.<https://aws.amazon.com/sqs/features/>QUESTION 108A company is migrating a distributed application to AWS. The application serves variable workloads. The legacy platform consists of a primary server that coordinates jobs across multiple compute nodes. The company wants to modernize the application with a solution that maximizes resiliency and scalability.How should a solutions architect design the architecture to meet these requirements?A. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a destination for the jobs.Implement the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group. Configure EC2 Auto Scaling to use scheduled scaling.B. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a destination for the jobs.Implement the compute nodes with Amazon EC2 Instances that are managed in an Auto Scaling group.Configure EC2 Auto Scaling based on the size of the queue.C. Implement the primary server and the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group. Configure AWS CloudTrail as a destination for the logs .Configure EC2 Auto Scaling based on the load on the primary server.D. Implement the primary server and the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group. Configure Amazon EventBridge (Amazon CloudWatch Events) as a destination for the jobs.Configure EC2 Auto Scaling based on the load on the compute nodes.Answer: BQUESTION 109A company is running an SMB file server in its data center. The file server stores large files that are accessed frequently for the first few days after the files are created. After 7 days the files are rarely accessed.The total data size is increasing and is close to the company's total storage capacity. A solutions architect must increase the company's available storage space without losing low-latency access to the most recently accessed files. The solutions architect must also provide file lifecycle management to avoid future storage issues.Which solution will meet these requirements?A. Use AWS DataSync to copy data that is older than 7 days from the SMB file server to AWS.B. Create an Amazon S3 File Gateway to extend the company's storage space. Create an S3 Lifecycle policy to transition the data to S3 Glacier Deep Archive after 7 days.C. Create an Amazon FSx for Windows File Server file system to extend the company's storage space.D. Install a utility on each user's computer to access Amazon S3. Create an S3 Lifecycle policy to transition the data to S3 Glacier Flexible Retrieval after 7 days.Answer: DQUESTION 110A company is building an ecommerce web application on AWS. The application sends information about new orders to an Amazon API Gateway REST API to process. The company wants to ensure that orders are processed in the order that they are received.Which solution will meet these requirements?A. Use an API Gateway integration to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic when the application receives an order. Subscribe an AWS Lambda function to the topic to perform processing.B. Use an API Gateway integration to send a message to an Amazon Simple Queue Service (Amazon SQS) FIFO queue when the application receives an order. Configure the SQS FIFO queue to invoke an AWS Lambda function for processing.C. Use an API Gateway authorizer to block any requests while the application processes an order.D. Use an API Gateway integration to send a message to an Amazon Simple Queue Service (Amazon SQS) standard queue when the application receives an order. Configure the SQS standard queue to invoke an AWS Lambda function for processing.Answer: AQUESTION 111A company has an application that runs on Amazon EC2 instances and uses an Amazon Aurora database. The EC2 instances connect to the database by using user names and passwords that are stored locally in a file. The company wants to minimize the operational overhead of credential management.What should a solutions architect do to accomplish this goal?A. Use AWS Secrets Manager. Turn on automatic rotation.B. Use AWS Systems Manager Parameter Store. Turn on automatic rotation.C. Create an Amazon S3 bucket to store objects that are encrypted with an AWS Key Management Service (AWS KMS) encryption

key. Migrate the credential file to the S3 bucket. Point the application to the S3 bucket.D. Create an encrypted Amazon Elastic Block Store (Amazon EBS) volume or each EC2 instance. Attach the new EBS volume to each EC2 instance. Migrate the credential file to the new EBS volume. Point the application to the new EBS volume. Answer: BQUESTION 112A global company hosts its web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The web application has static data and dynamic data. The company stores its static data in an Amazon S3 bucket. The company wants to improve performance and reduce latency for the static data and dynamic data. The company is using its own domain name registered with Amazon Route 53. What should a solutions architect do to meet these requirements?A. Create an Amazon CloudFront distribution that has the S3 bucket and the ALB as origins. Configure Route 53 to route traffic to the CloudFront distribution.B. Create an Amazon CloudFront distribution that has the ALB as an origin. Create an AWS Global Accelerator standard accelerator that has the S3 bucket as an endpoint. Configure Route 53 to route traffic to the CloudFront distribution.C. Create an Amazon CloudFront distribution that has the S3 bucket as an origin. Create an AWS Global Accelerator standard accelerator that has the ALB and the CloudFront distribution as endpoints. Create a custom domain name that points to the accelerator DNS name. Use the custom domain name as an endpoint for the web application.D. Create an Amazon CloudFront distribution that has the ALB as an origin. Create an AWS Global Accelerator standard accelerator that has the S3 bucket as an endpoint. Create two domain names. Point one domain name to the CloudFront DNS name for dynamic content. Point the other domain name to the accelerator DNS name for static content. Use the domain names as endpoints for the web application. Answer: AExplanation:

<https://stackoverflow.com/questions/52704816/how-to-properly-disable-cloudfront-caching-for-api-requests>QUESTION 113A company performs monthly maintenance on its AWS infrastructure. During these maintenance activities, the company needs to rotate the credentials for its Amazon RDS for MySQL databases across multiple AWS Regions. Which solution will meet these requirements with the LEAST operational overhead?A. Store the credentials as secrets in AWS Secrets Manager. Use multi-Region secret replication for the required Regions. Configure Secrets Manager to rotate the secrets on a schedule.B. Store the credentials as secrets in AWS Systems Manager by creating a secure string parameter. Use multi-Region secret replication for the required Regions. Configure Systems Manager to rotate the secrets on a schedule.C. Store the credentials in an Amazon S3 bucket that has server-side encryption (SSE) enabled. Use Amazon EventBridge (Amazon CloudWatch Events) to invoke an AWS Lambda function to rotate the credentials.D. Encrypt the credentials as secrets by using AWS Key Management Service (AWS KMS) multi-Region customer managed keys. Store the secrets in an Amazon DynamoDB global table. Use an AWS Lambda function to retrieve the secrets from DynamoDB. Use the RDS API to rotate the secrets. Answer: DQUESTION 114A company is planning to run a group of Amazon EC2 instances that connect to an Amazon Aurora database. The company has built an AWS CloudFormation template to deploy the EC2 instances and the Aurora DB cluster. The company wants to allow the instances to authenticate to the database in a secure way. The company does not want to maintain static database credentials. Which solution meets these requirements with the LEAST operational effort?A. Create a database user with a user name and password. Add parameters for the database user name and password to the CloudFormation template. Pass the parameters to the EC2 instances when the instances are launched.B. Create a database user with a user name and password. Store the user name and password in AWS Systems Manager Parameter Store. Configure the EC2 instances to retrieve the database credentials from Parameter Store.C. Configure the DB cluster to use IAM database authentication. Create a database user to use with IAM authentication. Associate a role with the EC2 instances to allow applications on the instances to access the database.D. Configure the DB cluster to use IAM database authentication with an IAM user. Create a database user that has a name that matches the IAM user. Associate the IAM user with the EC2 instances to allow applications on the instances to access the database. Answer: AExplanation: Finally, you need a way to instruct CloudFormation to complete stack creation only after all the services (such as Apache and MySQL) are running and not after all the stack resources are created. In other words, if you use the template from the earlier section to launch a stack, CloudFormation sets the status of the stack as CREATE\_COMPLETE after it successfully creates all the resources. However, if one or more services failed to start, CloudFormation still sets the stack status as CREATE\_COMPLETE. To prevent the status from changing to CREATE\_COMPLETE until all the services have successfully started, you can add a CreationPolicy attribute to the instance. This attribute puts the instance's status in CREATE\_IN\_PROGRESS until CloudFormation receives the required number of success signals or the timeout period is exceeded, so you can control when the instance has been successfully created. Reference:

[https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/deploying\\_applications.html](https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/deploying_applications.html)QUESTION 115A solutions architect is designing a shared storage solution for a web application that is deployed across multiple Availability Zones. The web application runs on Amazon EC2 instances that are in an Auto Scaling group. The company plans to make frequent changes to the content. The solution must have strong consistency in returning the new content as soon as the changes occur. Which solutions meet these requirements? (Select TWO.)A. Use AWS Storage Gateway Volume Gateway Internet Small Computer Systems Interface

(iSCSI) block storage that is mounted to the individual EC2 instances.B. Create an Amazon Elastic File System (Amazon EFS) file system. Mount the EFS file system on the individual EC2 instances.C. Create a shared Amazon Elastic Block Store (Amazon EBS) volume. Mount the EBS volume on the individual EC2 instances.D. Use AWS DataSync to perform continuous synchronization of data between EC2 hosts in the Auto Scaling group.E. Create an Amazon S3 bucket to store the web content. Set the metadata for the Cache-Control header to no-cache. Use Amazon CloudFront to deliver the content.

Answer: ABExplanation: Reference:<https://docs.aws.amazon.com/storagegateway/latest/userguide/WhatIsStorageGateway.html>

<https://docs.aws.amazon.com/efs/latest/ug/how-it-works.html> In this example, the EC2 instance in the us-west-2c Availability Zone will pay EC2 data access charges for accessing a mount target in a different Availability Zone. Creating this setup works as follows:

1. Create your Amazon EC2 resources and launch your Amazon EC2 instance. For more information about Amazon EC2, see Amazon EC2.
2. Create your Amazon EFS file system with One Zone storage.
3. Connect to each of your Amazon EC2 instances, and mount the Amazon EFS file system using the same mount target for each instance.

QUESTION 116A company that operates a web application on premises is preparing to launch a newer version of the application on AWS. The company needs to route requests to either the AWS-hosted or the on-premises-hosted application based on the URL query string. The on-premises application is not available from the internet, and a VPN connection is established between Amazon VPC and the company's data center. The company wants to use an Application Load Balancer (ALB) for this launch. Which solution meets these requirements?A. Use two ALBs: one for on-premises and one for the AWS resource. Add hosts to each target group of each ALB. Route with Amazon Route 53 based on the URL query string.B. Use two ALBs: one for on-premises and one for the AWS resource. Add hosts to the target group of each ALB. Create a software router on an EC2 instance based on the URL query string.C. Use one ALB with two target groups: one for the AWS resource and one for on-premises. Add hosts to each target group of the ALB. Configure listener rules based on the URL query string.D. Use one ALB with two AWS Auto Scaling groups: one for the AWS resource and one for on-premises. Add hosts to each Auto Scaling group. Route with Amazon Route 53 based on the URL query string.

Answer: C

Explanation: <https://aws.amazon.com/blogs/aws/new-advanced-request-routing-for-aws-application-load-balancers/> The host-based routing feature allows you to write rules that use the Host header to route traffic to the desired target group. Today we are extending and generalizing this feature, giving you the ability to write rules (and route traffic) based on standard and custom HTTP headers and methods, the query string, and the source IP address.

QUESTION 117A company wants to move from many standalone AWS accounts to a consolidated, multi-account architecture. The company plans to create many new AWS accounts for different business units. The company needs to authenticate access to these AWS accounts by using a centralized corporate directory service. Which combination of actions should a solutions architect recommend to meet these requirements? (Select TWO)

- A. Create a new organization in AWS Organizations with all features turned on. Create the new AWS accounts in the organization.
- B. Set up an Amazon Cognito identity pool. Configure AWS Single Sign-On to accept Amazon Cognito authentication.
- C. Configure a service control policy (SCP) to manage the AWS accounts. Add AWS Single Sign-On to AWS Directory Service.
- D. Create a new organization in AWS Organizations. Configure the organization's authentication mechanism to use AWS Directory Service directly.
- E. Set up AWS Single Sign-On (AWS SSO) in the organization. Configure AWS SSO and integrate it with the company's corporate directory service.

Answer: BCExplanation: SCPs affect only IAM users and roles that are managed by accounts that are part of the organization. SCPs don't affect resource-based policies directly. They also don't affect users or roles from accounts outside the organization. For example, consider an Amazon S3 bucket that's owned by account A in an organization. The bucket policy (a resource-based policy) grants access to users from account B outside the organization. Account A has an SCP attached. That SCP doesn't apply to those outside users in account B. The SCP applies only to users that are managed by account A in the organization.

An SCP restricts permissions for IAM users and roles in member accounts, including the member account's root user. Any account has only those permissions permitted by every parent above it. If a permission is blocked at any level above the account, either implicitly (by not being included in an Allow policy statement) or explicitly (by being included in a Deny policy statement), a user or role in the affected account can't use that permission, even if the account administrator attaches the AdministratorAccess IAM policy with /\* permissions to the user.

Reference: <https://aws.amazon.com/cognito/>

[https://docs.aws.amazon.com/organizations/latest/userguide/orgs\\_manage\\_policies\\_scps.html](https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html) QUESTION 118An entertainment company is using Amazon DynamoDB to store media metadata. The application is read intensive and experiencing delays. The company does not have staff to handle additional operational overhead and needs to improve the performance efficiency of DynamoDB without reconfiguring the application. What should a solutions architect recommend to meet this requirement?A. Use Amazon ElastiCache for Redis.

B. Use Amazon DynamoDB Accelerate (DAX).C. Replicate data by using DynamoDB global tables.D. Use Amazon ElastiCache for Memcached with Auto Discovery enabled.

Answer: BExplanation: Though DynamoDB offers consistent single-digit-millisecond latency, DynamoDB + DAX takes performance to the next level with response times in

microseconds for millions of requests per second for read-heavy workloads. With DAX, your applications remain fast and responsive, even when a popular event or news story drives unprecedented request volumes your way. No tuning required.

QUESTION 119 A company has an application that provides marketing services to stores. The services are based on previous purchases by store customers. The stores upload transaction data to the company through SFTP, and the data is processed and analyzed to generate new marketing offers. Some of the files can exceed 200 GB in size. Recently, the company discovered that some of the stores have uploaded files that contain personally identifiable information (PII) that should not have been included. The company wants administrators to be alerted if PII is shared again. The company also wants to automate remediation. What should a solutions architect do to meet these requirements with the LEAST development effort?

A. Use an Amazon S3 bucket as a secure transfer point. Use Amazon Inspector to scan the objects in the bucket. If objects contain PII, trigger an S3 Lifecycle policy to remove the objects that contain PII.

B. Use an Amazon S3 bucket as a secure transfer point. Use Amazon Macie to scan the objects in the bucket. If objects contain PII, use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects that contain PII.

C. Implement custom scanning algorithms in an AWS Lambda function. Trigger the function when objects are loaded into the bucket. If objects contain PII, use Amazon Simple Email Service (Amazon SES) to trigger a notification to the administrators and trigger an S3 Lifecycle policy to remove the objects that contain PII.

D. Implement custom scanning algorithms in an AWS Lambda function. Trigger the function when objects are loaded into the bucket. If objects contain PII, use Amazon Simple Email Service (Amazon SES) to trigger a notification to the administrators to remove the objects that contain PII.

Answer: B

QUESTION 120 A company needs guaranteed Amazon EC2 capacity in three specific Availability Zones in a specific AWS Region for an upcoming event that will last 1 week. What should the company do to guarantee the EC2 capacity?

A. Purchase Reserved instances that specify the Region needed.

B. Create an On Demand Capacity Reservation that specifies the Region needed.

C. Purchase Reserved instances that specify the Region and three Availability Zones needed.

D. Create an On-Demand Capacity Reservation that specifies the Region and three Availability Zones needed.

Explanation: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-capacity-reservations.html>

"When you create a Capacity Reservation, you specify the Availability Zone in which to reserve the capacity."

QUESTION 121 A company's website uses an Amazon EC2 instance store for its catalog of items. The company wants to make sure that the catalog is highly available and that the catalog is stored in a durable location. What should a solutions architect do to meet these requirements?

A. Move the catalog to Amazon ElastiCache for Redis.

B. Deploy a larger EC2 instance with a larger instance store.

C. Move the catalog from the instance store to Amazon S3 Glacier Deep Archive.

D. Move the catalog to an Amazon Elastic File System (Amazon EFS) file system.

Answer: A

QUESTION 122 A company stores call transcript files on a monthly basis. Users access the files randomly within 1 year of the call, but users access the files infrequently after 1 year. The company wants to optimize its solution by giving users the ability to query and retrieve files that are less than 1-year-old as quickly as possible. A delay in retrieving older files is acceptable. Which solution will meet these requirements MOST cost-effectively?

A. Store individual files with tags in Amazon S3 Glacier Instant Retrieval. Query the tags to retrieve the files from S3 Glacier Instant Retrieval.

B. Store individual files in Amazon S3 Intelligent-Tiering. Use S3 Lifecycle policies to move the files to S3 Glacier Flexible Retrieval after 1 year. Query and retrieve the files that are in Amazon S3 by using Amazon Athena.

C. Store individual files with tags in Amazon S3 Standard storage. Store search metadata for each archive in Amazon S3 Standard storage. Use S3 Lifecycle policies to move the files to S3 Glacier Instant Retrieval after 1 year. Query and retrieve the files by searching for metadata from Amazon S3.

D. Store individual files in Amazon S3 Standard storage. Use S3 Lifecycle policies to move the files to S3 Glacier Deep Archive after 1 year. Store search metadata in Amazon RDS. Query the files from Amazon RDS. Retrieve the files from S3 Glacier Deep Archive.

Answer: C

QUESTION 123 A company has a production workload that runs on 1,000 Amazon EC2 Linux instances. The workload is powered by third-party software. The company needs to patch the third-party software on all EC2 instances as quickly as possible to remediate a critical security vulnerability. What should a solutions architect do to meet these requirements?

A. Create an AWS Lambda function to apply the patch to all EC2 instances.

B. Configure AWS Systems Manager Patch Manager to apply the patch to all EC2 instances.

C. Schedule an AWS Systems Manager maintenance window to apply the patch to all EC2 instances.

D. Use AWS Systems Manager Run Command to run a custom command that applies the patch to all EC2 instances.

Answer: D

QUESTION 124 A company is developing an application that provides order shipping statistics for retrieval by a REST API. The company wants to extract the shipping statistics, organize the data into an easy-to-read HTML format, and send the report to several email addresses at the same time every morning. Which combination of steps should a solutions architect take to meet these requirements? (Choose two.)

A. Configure the application to send the data to Amazon Kinesis Data Firehose.

B. Use Amazon Simple Email Service (Amazon SES) to format the data and to send the report by email.

C. Create an Amazon EventBridge (Amazon CloudWatch Events) scheduled event that invokes an AWS Glue job to query

the application's API for the data.D. Create an Amazon EventBridge (Amazon CloudWatch Events) scheduled event that invokes an AWS Lambda function to query the application's API for the data.E. Store the application data in Amazon S3. Create an Amazon Simple Notification Service (Amazon SNS) topic as an S3 event destination to send the report byAnswer: DEQUESTION 125A company wants to migrate its on-premises application to AWS. The application produces output files that vary in size from tens of gigabytes to hundreds of terabytes The application data must be stored in a standard file system structure. The company wants a solution that scales automatically, is highly available, and requires minimum operational overhead.Which solution will meet these requirements?A. Migrate the application to run as containers on Amazon Elastic Container Service (Amazon ECS).Use Amazon S3 for storage.B. Migrate the application to run as containers on Amazon Elastic Kubernetes Service (Amazon EKS).Use Amazon Elastic Block Store (Amazon EBS) for storage.C. Migrate the application to Amazon EC2 instances in a Multi-AZ Auto Scaling group. Use Amazon Elastic File System (Amazon EFS) for storage.D. Migrate the application to Amazon EC2 instances in a Multi-AZ Auto Scaling group. Use Amazon Elastic Block Store (Amazon EBS) for storage.Answer: CResources From:1.2022 Latest Braindump2go SAA-C03 Exam Dumps (PDF & VCE) Free Share:<https://www.braindump2go.com/saa-c03.html>2.2022 Latest Braindump2go SAA-C03 PDF and SAA-C03 VCE Dumps Free Share:  
[https://drive.google.com/drive/folders/1PKc\\_AsNW5xtYjJaY4\\_oJFcTRLkRk9lPW?usp=sharing](https://drive.google.com/drive/folders/1PKc_AsNW5xtYjJaY4_oJFcTRLkRk9lPW?usp=sharing)3.2021 Free Braindump2go SAA-C03 Exam Questions Download:[https://www.braindump2go.com/free-online-pdf/SAA-C03-Dumps\(66-90\).pdf](https://www.braindump2go.com/free-online-pdf/SAA-C03-Dumps(66-90).pdf)  
[https://www.braindump2go.com/free-online-pdf/SAA-C03-PDF\(21-41\).pdf](https://www.braindump2go.com/free-online-pdf/SAA-C03-PDF(21-41).pdf)  
[https://www.braindump2go.com/free-online-pdf/SAA-C03-PDF-Dumps\(1-20\).pdf](https://www.braindump2go.com/free-online-pdf/SAA-C03-PDF-Dumps(1-20).pdf)  
[https://www.braindump2go.com/free-online-pdf/SAA-C03-VCE\(42-65\).pdf](https://www.braindump2go.com/free-online-pdf/SAA-C03-VCE(42-65).pdf)  
[https://www.braindump2go.com/free-online-pdf/SAA-C03-VCE-Dumps\(90-110\).pdf](https://www.braindump2go.com/free-online-pdf/SAA-C03-VCE-Dumps(90-110).pdf)Free Resources from Braindump2go,We Devoted to Helping You 100% Pass All Exams!