Pass2Lead
https://Pass2Lead.com

SAP-C02^{Q&As}

AWS Certified Solutions Architect - Professional

Pass Amazon SAP-C02 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.pass2lead.com/sap-c02.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon Official Exam Center

⚙ **Instant Download** After Purchase

⚙ **100% Money Back** Guarantee

⚙ **365 Days** Free Update

⚙ **800,000+** Satisfied Customers

QUESTION 1

A company runs a new application as a static website in Amazon S3. The company has deployed the application to a production AWS account and uses Amazon CloudFront to deliver the website. The website calls an Amazon API Gateway REST API. An AWS Lambda function backs each API method.

The company wants to create a CSV report every 2 weeks to show each API Lambda function\\'s recommended configured memory, recommended cost, and the price difference between current configurations and the recommendations. The company will store the reports in an S3 bucket.

Which solution will meet these requirements with the LEAST development time?

A. Create a Lambda function that extracts metrics data for each API Lambda function from Amazon CloudWatch Logs for the 2-week penod_ Collate the data into tabular format. Store the data as a _csvfile in an S3 bucket. Create an Amazon Eventaridge rule to schedule the Lambda function to run every 2 weeks.

B. Opt in to AWS Compute Optimizer. Create a Lambda function that calls the ExportLambdaFunctionRecommendatlons operation. Export the _csv file to an S3 bucket. Create an Amazon Eventaridge rule to schedule the Lambda function to run every 2 weeks.

C. Opt in to AWS Compute Optimizer. Set up enhanced infrastructure metrics. Within the Compute Optimizer console, schedule a job to export the Lambda recommendations to a _csvfile_ Store the file in an S3 bucket every 2 weeks.

D. Purchase the AWS Business Support plan for the production account. Opt in to AWS Compute Optimizer for AWS Trusted Advisor checks. In the Trusted Advisor console, schedule a job to export the cost optimization checks to a _csvfile_ Store the file in an S3 bucket every 2 weeks.

Correct Answer: B

https://docs.aws.amazon.com/compute-optimizer/latest/APIReference/API_ExportLambdaFunctionRecommendations.html

QUESTION 2

A company uses Amazon S3 to store files and images in a variety of storage classes. The company\\'s S3 costs have increased substantially during the past year.

A solutions architect needs to review data trends for the past 12 months and identity the appropriate storage class for the objects.

Which solution will meet these requirements?

A. Download AWS Cost and Usage Reports for the last 12 months of S3 usage. Review AWS Trusted Advisor recommendations for cost savings.

B. Use S3 storage class analysis. Import data trends into an Amazon QuickSight dashboard to analyze storage trends.

C. Use Amazon S3 Storage Lens. Upgrade the default dashboard to include advanced metrics for storage trends.

D. Use Access Analyzer for S3. Download the Access Analyzer for S3 report for the last 12 months. Import the csvfile to an Amazon QuickSight dashboard.

Correct Answer: B

![Pass2Lead](https://Pass2Lead.com)
Amazon S3 Storage Class Analysis allows you to analyze your object access patterns and automatically move data to the most cost-effective storage class, without any code changes. This feature can automatically transition objects to the appropriate storage class based on the access patterns. You can also use the S3 storage class analysis feature to import data trends into an Amazon QuickSight dashboard and analyze storage trends for the past 12 months, to identify the appropriate storage class for the objects. AWS Cost and Usage Reports provide detailed information about your AWS usage and costs, but it does not provide specific storage class analysis. AWS Trusted Advisor can provide cost savings recommendations but it does not provide specific storage class analysis. Amazon S3 Storage Lens, provides detailed metrics on S3 storage usage, but it does not provide specific storage class analysis. Access Analyzer for S3, analyzes access control lists (ACLs) and resource-based policies in S3 buckets to identify S3 bucket access that is open to the public, but it does not provide specific storage class analysis. References: Amazon S3 Storage Class Analysis AWS Cost and Usage Reports AWS Trusted Advisor Amazon S3 Storage Lens Access Analyzer for S3

QUESTION 3

A company migrated an application to the AWS Cloud. The application runs on two Amazon EC2 instances behind an Application Load Balancer (ALB). Application data is stored in a MySQL database that runs on an additional EC2 instance. The application\\'s use of the database is read-heavy.

The loads static content from Amazon Elastic Block Store (Amazon EBS) volumes that are attached to each EC2 instance. The static content is updated frequently and must be copied to each EBS volume.

The load on the application changes throughout the day. During peak hours, the application cannot handle all the incoming requests. Trace data shows that the database cannot handle the read load during peak hours.

Which solution will improve the reliability of the application?

A. Migrate the application to a set of AWS Lambda functions. Set the Lambda functions as targets for the ALB. Create a new single EBS volume for the static content. Configure the Lambda functions to read from the new EBS volume. Migrate the database to an Amazon RDS for MySQL Multi-AZ DB cluster.

B. Migrate the application to a set of AWS Step Functions state machines. Set the state machines as targets for the ALB. Create an Amazon Elastic File System (Amazon EFS) file system for the static content. Configure the state machines to read from the EFS file system. Migrate the database to Amazon Aurora MySQL Serverless v2 with a reader DB instance.

C. Containerize the application. Migrate the application to an Amazon Elastic Container Service (Amazon ECS) Cluster. Use the AWS Fargate launch type for the tasks that host the application. Create a new single EBS volume the static content. Mount the new EBS volume on the ECS duster. Configure AWS Application Auto Scaling on ECS cluster. Set the ECS service as a target for the ALB. Migrate the database to an Amazon RDS for MySOL Multi-AZ DB cluster.

D. Containerize the application. Migrate the application to an Amazon Elastic Container Service (Amazon ECS) cluster. Use the AWS Fargate launch type for the tasks that host the application. Create an Amazon Elastic File System (Amazon EFS) file system for the static content. Mount the EFS file system to each container. Configure AWS Application Auto Scaling on the ECS cluster Set the ECS service as a target for the ALB. Migrate the database to Amazon Aurora MySQL Serverless v2 with a reader DB instance.

Correct Answer: D

This solution will improve the reliability of the application by addressing the issues of scalability, availability, and performance. Containerizing the application will make it easier to deploy and manage on AWS. Migrating the application to an Amazon ECS cluster will allow the application to run on a fully managed container orchestration service. Using the AWS Fargate launch type for the tasks that host the application will enable the application to run on serverless compute engines that are automatically provisioned and scaled by AWS. Creating an Amazon EFS file system for the static content will provide a scalable and shared storage solution that can be accessed by multiple containers. Mounting the EFS file system to each container will eliminate the need to copy the static content to each EBS volume and ensure that

the content is always up to date. Configuring AWS Application Auto Scaling on the ECS cluster will enable the application to scale up and down based on demand or a predefined schedule. Setting the ECS service as a target for the ALB will distribute the incoming requests across multiple tasks in the ECS cluster and improve the availability and fault tolerance of the application. Migrating the database to Amazon Aurora MySQL Serverless v2 with a reader DB instance will provide a fully managed, compatible, and scalable relational database service that can handle high throughput and concurrent connections. Using a reader DB instance will offload some of the read load from the primary DB instance and improve the performance of the database.

QUESTION 4

A gaming company created a game leaderboard by using a Multi-AZ deployment of an Amazon RDS database. The number of users is growing, and the queries to get individual player rankings are getting slower over time. The company expects a surge in users for an upcoming version and wants to optimize the design for scalability and performance.

Which solution will meet these requirements?

A. Migrate the database to Amazon DynamoDB. Store the leader different tables. Use Apache HiveQL JOIN statements to build the leaderboard

B. Keep the leaderboard data in the RDS DB instance. Provision a Multi-AZ deployment of an Amazon ElastiCache for Redis cluster.

C. Stream the leaderboard data by using Amazon Kinesis Data Firehose with an Amazon S3 bucket as the destination. Query the S3 bucket by using Amazon Athena for the leaderboard.

D. Add a read-only replica to the RDS DB instance. Add an RDS Proxy database proxy.

Correct Answer: C

QUESTION 5

A media storage application uploads user photos to Amazon S3 for processing by AWS Lambda functions. Application state is stored in Amazon DynamoDB tables. Users are reporting that some uploaded photos are not being processed properly. The application developers trace the logs and find that Lambda is experiencing photo processing issues when thousands of users upload photos simultaneously. The issues are the result of Lambda concurrency limits and the performance of DynamoDB when data is saved.

Which combination of actions should a solutions architect take to increase the performance and reliability of the application? (Select TWO.)

A. Evaluate and adjust the RCUs tor the DynamoDB tables.

B. Evaluate and adjust the WCUs for the DynamoDB tables.

C. Add an Amazon ElastiCache layer to increase the performance of Lambda functions.

D. Add an Amazon Simple Queue Service (Amazon SQS) queue and reprocessing logic between Amazon S3 and the Lambda functions.

E. Use S3 Transfer Acceleration to provide lower latency to users.

Correct Answer: BD

![Pass2Lead](https://Pass2Lead.com)
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.html#Ho
wItWorks.requests https://aws.amazon.com/blogs/compute/robust-serverless-application-design-with-aws-lambda-dlq/c

QUESTION 6

A large mobile gaming company has successfully migrated all of its on-premises infrastructure to the AWS Cloud. A
solutions architect is reviewing the environment to ensure that it was built according to the design and that it is running
in alignment with the Well-Architected Framework.

While reviewing previous monthly costs in Cost Explorer, the solutions architect notices that the creation and
subsequent termination of several large instance types account for a high proportion of the costs. The solutions architect
finds out that the company\\'s developers are launching new Amazon EC2 instances as part of their testing and that the
developers are not using the appropriate instance types.

The solutions architect must implement a control mechanism to limit the instance types that only the developers can
launch.

Which solution will meet these requirements?

A. Create a desired-instance-type managed rule in AWS Config. Configure the rule with the instance types that are
allowed. Attach the rule to an event to run each time a new EC2 instance is launched.

B. In the EC2 console, create a launch template that specifies the instance types that are allowed. Assign the launch
template to the developers\\' IAM accounts.

C. Create a new IAM policy. Specify the instance types that are allowed. Attach the policy to an IAM group that contains
the IAM accounts for the developers

D. Use EC2 Image Builder to create an image pipeline for the developers and assist them in the creation of a golden
image.

Correct Answer: C

This is doable with IAM policy creation to restrict users to specific instance types. Found the below
article.https://blog.vizuri.com/limiting-allowed-aws-instance-type-with-iam-policy

QUESTION 7

A financial services company receives a regular data feed from its credit card servicing partner Approximately 5.000
records are sent every 15 minutes in plaintext, delivered over HTTPS directly into an Amazon S3 bucket with server-
side encryption. This feed contains sensitive credit card primary account number (PAN) data The company needs to
automatically mask the PAN before sending the data to another S3 bucket for additional internal processing. The
company also needs to remove and merge specific fields, and then transform the record into JSON format Additionally,
extra feeds are likely to be added in the future, so any design needs to be easily expandable.

Which solutions will meet these requirements?

A. Trigger an AWS Lambda function on file delivery that extracts each record and writes it to an Amazon SQS queue.
Trigger another Lambda function when new messages arrive in the SQS queue to process the records, writing the
results to a temporary location in Amazon S3. Trigger a final Lambda function once the SQS queue is empty to
transform the records into JSON format and send the results to another S3 bucket for internal processing.

B. Trigger an AWS Lambda function on file delivery that extracts each record and writes it to an Amazon SQS queue.

Configure an AWS Fargate container application to automatically scale to a single instance when the SQS queue contains messages. Have the application process each record, and transform the record into JSON format. When the queue is empty, send the results to another S3 bucket for internal processing and scale down the AWS Fargate instance.

C. Create an AWS Glue crawler and custom classifier based on the data feed formats and build a table definition to match. Trigger an AWS Lambda function on file delivery to start an AWS Glue ETL job to transform the entire record according to the processing and transformation requirements. Define the output format as JSON. Once complete, have the ETL job send the results to another S3 bucket for internal processing.

D. Create an AWS Glue crawler and custom classifier based upon the data feed formats and build a table definition to match. Perform an Amazon Athena query on file delivery to start an Amazon EMR ETL job to transform the entire record according to the processing and transformation requirements. Define the output format as JSON. Once complete, send the results to another S3 bucket for internal processing and scale down the EMR cluster.

Correct Answer: C

You can use a Glue crawler to populate the AWS Glue Data Catalog with tables. The Lambda function can be triggered using S3 event notifications when object create events occur. The Lambda function will then trigger the Glue ETL job to transform the records masking the sensitive data and modifying the output format to JSON. This solution meets all requirements.

Create an AWS Glue crawler and custom classifier based on the data feed formats and build a table definition to match. Trigger an AWS Lambda function on file delivery to start an AWS Glue ETL job to transform the entire record according to the processing and transformation requirements. Define the output format as JSON. Once complete, have the ETL job send the results to another S3 bucket for internal processing. https://docs.aws.amazon.com/glue/latest/dg/trigger-job.html https://d1.awsstatic.com/Products/product-name/diagrams/product-page-diagram_Glue_Event-driven-ETLPipelines.e24d59bb79a9e24cdba7f43ffd234ec0482a60e2.png

QUESTION 8

A company runs an application in the cloud that consists of a database and a website Users can post data to the website, have the data processed, and have the data sent back to them in an email. Data is stored in a MySQL database running on an Amazon EC2 instance The database is running in a VPC with two private subnets The website is running on Apache Tomcat in a single EC2 instance in a different VPC with one public subnet There is a single VPC peering connection between the database and website VPC.

The website has suffered several outages during the last month due to high traffic Which actions should a solutions architect take to increase the reliability of the application? (Select THREE)

A. Place the Tomcat server in an Auto Scaling group with multiple EC2 instances behind an Application Load Balancer

B. Provision an additional VPC peering connection

C. Migrate the MySQL database to Amazon Aurora with one Aurora Replica

D. Provision two NAT gateways in the database VPC

E. Move the Tomcat server to the database VPC

F. Create an additional public subnet in a different Availability Zone in the website VPC

Correct Answer: ACF

QUESTION 9

A global manufacturing company plans to migrate the majority of its applications to AWS. However, the company is concerned about applications that need to remain within a specific country or in the company\'s central on-premises data center because of data regulatory requirements or requirements for latency of single-digit milliseconds. The company also is concerned about the applications that it hosts in some of its factory sites, where limited network infrastructure exists.

The company wants a consistent developer experience so that its developers can build applications once and deploy on premises, in the cloud, or in a hybrid architecture.

The developers must be able to use the same tools, APIs, and services that are familiar to them.

Which solution will provide a consistent hybrid experience to meet these requirements?

A. Migrate all applications to the closest AWS Region that is compliant. Set up an AWS Direct Connect connection between the central on-premises data center and AWS. Deploy a Direct Connect gateway.

B. Use AWS Snowball Edge Storage Optimized devices for the applications that have data regulatory requirements or requirements for latency of single-digit milliseconds. Retain the devices on premises. Deploy AWS Wavelength to host the workloads in the factory sites.

C. Install AWS Outposts for the applications that have data regulatory requirements or requirements for latency of single-digit milliseconds. Use AWS Snowball Edge Compute Optimized devices to host the workloads in the factory sites.

D. Migrate the applications that have data regulatory requirements or requirements for latency of single-digit milliseconds to an AWS Local Zone. Deploy AWS Wavelength to host the workloads in the factory sites.

Correct Answer: C

Installing AWS Outposts for the applications that have data regulatory requirements or requirements for latency of single-digit milliseconds will provide a fully managed service that extends AWS infrastructure, services, APIs, and tools to customer premises1. AWS Outposts allows customers to run some AWS services locally and connect to a broad range of services available in the local AWS Region1. Using AWS Snowball Edge Compute Optimized devices to host the workloads in the factory sites will provide local compute and storage resources for locations with limited network infrastructure2. AWS Snowball Edge devices can run Amazon EC2 instances and AWS Lambda functions locally and sync data with AWS when network connectivity is available2.

QUESTION 10

A company is running an application on Amazon EC2 instances in the AWS Cloud. The application is using a MongoDB database with a replica set as its data tier. The MongoDB database is installed on systems in the company\'s on-premises data center and is accessible through an AWS Direct Connect connection to the data center environment.

A solutions architect must migrate the on-premises MongoDB database to Amazon DocumentDB (with MongoDB compatibility).

Which strategy should the solutions architect choose to perform this migration?

A. Create a fleet of EC2 instances. Install MongoDB Community Edition on the EC2 instances, and create a database. Configure continuous synchronous replication with the database that is running in the on-premises data center.

B. Create an AWS Database Migration Service (AWS DMS) replication instance. Create a source endpoint for the on-premises MongoDB database by using change data capture (CDC). Create a target endpoint for the Amazon DocumentDB database. Create and run a DMS migration task.

![Pass2Lead](https://Pass2Lead.com)

C. Create a data migration pipeline by using AWS Data Pipeline. Define data nodes for the on-premises MongoDB database and the Amazon DocumentDB database. Create a scheduled task to run the data pipeline.

D. Create a source endpoint for the on-premises MongoDB database by using AWS Glue crawlers. Configure continuous asynchronous replication between the MongoDB database and the Amazon DocumentDB database.

Correct Answer: B

https://aws.amazon.com/getting-started/hands-on/move-to-managed/migrate-mongodb-to-documentdb/

QUESTION 11

A solutions architect is building a web application that uses an Amazon RDS for PostgreSQL DB instance The DB instance is expected to receive many more reads than writes. The solutions architect needs to ensure that the large amount of read traffic can be accommodated and that the DB instance is highly available.

Which steps should the solutions architect take to meet these requirements? (Select THREE)

A. Create multiple read replicas and put them into an Auto Scaling group.

B. Create multiple read replicas in different Availability Zones.

C. Create an Amazon Route 53 hosted zone and a record set for each read replica with a TTL and a weighted routing policy.

D. Create an Application Load Balancer (ALB) and put the read replicas behind the ALB.

E. Configure an Amazon CloudWatch alarm to detect a failed read replica. Set the alarm to directly invoke an AWS Lambda function to delete its Route 53 record set.

F. Configure an Amazon Route 53 health check for each read replica using its endpoint

Correct Answer: BCF

https://aws.amazon.com/premiumsupport/knowledge-center/requests-rds- read-replicas/ You can use Amazon Route 53 weighted record sets to distribute requests across your read replicas. Within a Route 53 hosted zone, create individual record sets for each DNS endpoint associated with your read replicas and give them the same weight. Then, direct requests to the endpoint of the record set. You can incorporate Route 53 health checks to be sure that Route 53 directs traffic away from unavailable read replicas

---

QUESTION 12

During 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

![Pass2Lead](https://Pass2Lead.com)
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.

Correct Answer: D

CodeCommit may use S3 on the back end (and it also uses DynamoDB on the back end) but I don\\'t think they\\'re stored in buckets that you can see or point Macie to. In fact, there are even solutions out there describing how to copy your repo from CodeCommit into S3 to back it up: https://docs.aws.amazon.com/prescriptive- guidance/latest/patterns/a utomate-event-driven-backups-from-codecommit-to-amazon-s3-using-codebuild-and-cloudwatch-events.html

QUESTION 13

A financial services company runs a complex, multi-tier application on Amazon EC2 instances and AWS Lambda functions. The application stores temporary data in Amazon S3. The S3 objects are valid for only 45 minutes and are deleted after 24 hours.

The company deploys each version of the application by launching an AWS CloudFormation stack. The stack creates all resources that are required to run the application. When the company deploys and validates a new application version, the company deletes the CloudFormation stack of the old version.

The company recently tried to delete the CloudFormation stack of an old application version, but the operation failed. An analysis shows that CloudFormation failed to delete an existing S3 bucket. A solutions architect needs to resolve this issue without making major changes to the application\\'s architecture.

Which solution meets these requirements?

A. Implement a Lambda function that deletes all files from a given S3 bucket. Integrate this Lambda function as a custom resource into the CloudFormation stack. Ensure that the custom resource has a DependsOn attribute that points to the S3 bucket\\'s resource.

B. Modify the CloudFormation template to provision an Amazon Elastic File System (Amazon EFS) file system to store the temporary files there instead of in Amazon S3. Configure the Lambda functions to run in the same VPC as the file system. Mount the file system to the EC2 instances and Lambda functions.

C. Modify the CloudFormation stack to create an S3 Lifecycle rule that expires all objects 45 minutes after creation. Add a DependsOn attribute that points to the S3 bucket\\'s resource.

D. Modify the CloudFormation stack to attach a DeletionPolicy attribute with a value of Delete to the S3 bucket.

Correct Answer: D

This option allows the solutions architect to use a DeletionPolicy attribute to specify how AWS CloudFormation handles the deletion of an S3 bucket when the stack is deleted1. By setting the value of Delete, the solutions architect can instruct

CloudFormation to delete the bucket and all of its contents1. This option does not require any major changes to the application\\'s architecture or any additional resources.

References:

Deletion policies

QUESTION 14

A company has multiple business units that each have separate accounts on AWS. Each business unit manages its own network with several VPCs that have CIDR ranges that overlap. The company\\'s marketing team has created a new internal application and wants to make the application accessible to all the other business units. The solution must use private IP addresses only.

Which solution will meet these requirements with the LEAST operational overhead?

A. Instruct each business unit to add a unique secondary CIDR range to the business unit\\'s VPC. Peer the VPCs and use a private NAT gateway in the secondary range to route traffic to the marketing team.

B. Create an Amazon EC2 instance to serve as a virtual appliance in the marketing account\\'s VPC. Create an AWS Site-to-Site VPN connection between the marketing team and each business unit\\'s VPC. Perform NAT where necessary.

C. Create an AWS PrivateLink endpoint service to share the marketing application. Grant permission to specific AWS accounts to connect to the service. Create interface VPC endpoints in other accounts to access the application by using private IP addresses.

D. Create a Network Load Balancer (NLB) in front of the marketing application in a private subnet. Create an API Gateway API. Use the Amazon API Gateway private integration to connect the API to the NLB. Activate IAM authorization for the API. Grant access to the accounts of the other business units.

Correct Answer: C

With AWS PrivateLink, the marketing team can create an endpoint service to share their internal application with other accounts securely using private IP addresses. They can grant permission to specific AWS accounts to connect to the service and create interface VPC endpoints in the other accounts to access the application by using private IP addresses. This option does not require any changes to the network of the other business units, and it does not require peering or NATing. This solution is both scalable and secure. https://aws.amazon.com/blogs/networking-and-content-delivery/connecting-networks-with-overlapping-ip-ranges/

QUESTION 15

A large education company recently introduced Amazon Workspaces to provide access to internal applications across multiple universities. The company is storing user proxies on an Amazon FSx for Windows File Server tile system. The Me system is configured with a DNS alias and is connected to a self-managed Active Directory As more users begin to use the Workspaces login time increases to unacceptable levels

An investigation reveals a degradation in performance of the file system. The company created the file system on HDD storage with a throughput of 16 MBps A solutions architect must improve the performance of the file system during a defined maintenance window

What should the solutions architect do to meet these requirements with the LEAST administrative effort?

A. Use AWS Backup to create a point-in-time backup of the file system Restore the backup to a new FSx for Windows File Server file system Select SSD as the storage type Select 32 MBps as the throughput capacity When the backup and restore process is completed adjust the DNS alias accordingly Delete the original file system B. Disconnect users from the file system In the Amazon FSx console, update the throughput capacity to 32 MBps Update the storage type to SSD Reconnect users to the file system

C. Deploy an AWS DataSync agent onto a new Amazon EC2 instance. Create a task Configure the existing file system

![Pass2Lead](https://Pass2Lead.com)
as the source location Configure a new FSx for Windows File Server file system with SSD storage and 32 MBps of throughput as the target location Schedule the task When the task is completed adjust the DNS alias accordingly Delete the original file system.

D. Enable shadow copies on the existing file system by using a Windows PowerShell command Schedule the shadow copy job to create a point-in-time backup of the file system Choose to restore previous versions Create a new FSx for Windows File Server file system with SSD storage and 32 MBps of throughput When the copy job is completed, adjust the DNS alias Delete the original file system

Correct Answer: A

[Latest SAP-C02 Dumps](https://www.pass2lead.com/sap-c02.html)   [SAP-C02 Practice Test](https://www.pass2lead.com/sap-c02.html)   [SAP-C02 Study Guide](https://www.pass2lead.com/sap-c02.html)