



PAS-C01^{Q&As}

AWS Certified: SAP on AWS - Specialty exam

Pass Amazon PAS-C01 Exam with 100% Guarantee

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

<https://www.passapply.com/pas-c01.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 core business processes on SAP. The company plans to migrate its SAP workloads to AWS.

Which combination of prerequisite steps must the company take to receive integrated support for SAP on AWS? (Choose three.)

- A. Purchase an AWS Developer Support plan or an AWS Enterprise Support plan.
- B. Purchase an AWS Business Support plan or an AWS Enterprise Support plan.
- C. Enable Amazon CloudWatch detailed monitoring.
- D. Enable Amazon EC2 termination protection.
- E. Configure and run the AWS Data Provider for SAP agent.
- F. Use Reserved Instances for all Amazon EC2 instances that run SAP.

Correct Answer: ACE

QUESTION 2

A company is running its on-premises SAP ERP Central Component (SAP ECC) system on an Oracle database on Oracle Enterprise Linux. The database is 1 TB in size and uses 27,000 IOPS for its peak performance. Multiple SSD volumes are striped to store Oracle data files in separate sapdata directories to gain the required IOPS.

The company is planning to move this workload to AWS. The company chooses high I/O bandwidth instances with a Nitro hypervisor to host the target database instance. Downtime is not a constraint for the migration. The company needs an Amazon Elastic Block Store (Amazon EBS) storage layout that optimizes cost for the migration.

How should the company reorganize the Oracle data files to meet these requirements?

- A. Reorganize the Oracle data files into one 9 TB General Purpose SSD (gp2) EBS volume.
- B. Reorganize the Oracle data files into a striped volume of three 3 TB General Purpose SSD (gp2) EBS volumes.
- C. Reorganize the Oracle data files into one 1 TB General Purpose SSD (gp3) EBS volume with 27,000 provisioned IOPS.
- D. Reorganize the Oracle data files into ten 100 GB General Purpose SSD (gp3) EBS volumes.

Correct Answer: D

QUESTION 3

A company is planning to migrate its SAP Content Server from on premises to Amazon EC2 instances. The SAP



Content Server stores data in a MaxDB database. The on-premises servers run the SUSE Linux Enterprise Server operating system.

The company wants to assess the benefits of cloud deployment by performing a proof of concept. An SAP solutions architect needs to perform a rehosting of the SAP Content Server on AWS to provide highly available and resilient storage.

Which solutions will meet these requirements? (Choose two.)

- A. Configure Amazon Elastic File System (Amazon EFS) file systems for the MaxDB permanent storage. Install the nfs-utils package on the EC2 instances. Create the necessary mounts to attach the EFS file systems to the EC2 instances.
- B. Configure Amazon FSx for Lustre file systems for the MaxDB permanent storage. Create the necessary mounts to attach the FSx for Lustre file systems to the EC2 instances. Update the /etc/fstab file with the directory name, DNS name, and mount name.
- C. Configure General Purpose SSD (gp2 or gp3) or Provisioned IOPS SSD (io1 or io2) Amazon Elastic Block Store (Amazon EBS) volumes for the MaxDB permanent storage. Use the aws ec2 attach-volume AWS CLI command with device, volume ID, and instance ID to attach the mount to each EC2 instance.
- D. Configure Amazon S3 buckets for the MaxDB permanent storage. Create an IAM instance profile that specifies a role to grant access to Amazon S3. Attach the instance profile to the EC2 instances.
- E. Configure Amazon Elastic Container Service (Amazon ECS) volumes for the MaxDB permanent storage. Install the nfs-utils package on the EC2 instances. Create the necessary mounts to attach the ECS volumes to the EC2 instances.

Correct Answer: AC

QUESTION 4

A company is running an SAP ERP Central Component (SAP ECC) system on an SAP HANA database that is 10 TB in size. The company is receiving notifications about long-running database backups every day. The company uses AWS Backup Agent for SAP HANA (AWS Backup agent) on an Amazon EC2 instance to back up the database. An SAP NetWeaver administrator needs to troubleshoot the problem and propose a solution.

Which solution will help resolve this problem?

- A. Ensure that the AWS Backup agent is configured to send the backups to an Amazon S3 bucket over the internet. Ensure that the EC2 instance is configured to access the internet through a NAT gateway.
- B. Check the UploadChannelSize parameter for the AWS Backup agent. Increase this value in the aws-backup-agent-config.yaml configuration file based on the EC2 instance type and storage configurations.
- C. Check the MaximumConcurrentFilesForRestore parameter for the AWS Backup agent. Increase the parameter from 5 to 10 by using the aws-backup-agent-config.yaml configuration file.
- D. Ensure that the backups are compressed if necessary. Configure the AWS Backup agent to compress the backups and send them to an Amazon S3 bucket.

Correct Answer: A

**QUESTION 5**

A company wants to improve the RPO and RTO for its SAP disaster recovery (DR) solution by running the DR solution on AWS. The company is running SAP ERP Central Component (SAP ECC) on SAP HANA. The company has set an RPO of 15 minutes and an RTO of 4 hours.

The production SAP HANA database is running on a physical appliance that has x86 architecture. The appliance has 1 TB of memory and the SAP HANA global allocation limit is set to 768 GB. The SAP application servers are running as VMs on VMware and they store data on an NFS file system. The company does not want to change any existing SAP HANA parameters that are related to data and log backup for its on-premises systems.

What should an SAP solutions architect do to meet the DR objectives MOST cost-effectively?

- A. For the SAP HANA database change the log backup frequency to 5 minutes. Move the data and log backups to Amazon S3 by using the AWS CLI or AWS DataSync. Launch the SAP HANA database. For the SAP application servers, export the VMs as AMIs by using the VM import/export feature from AWS. For NFS file shares /sapmnt and /usr/sap/trans/ establish real-time synchronization from DataSync to Amazon Elastic File System (Amazon EFS).
- B. For the SAP HANA database change the log backup frequency to 5 minutes. Move the data and log backups to Amazon S3 by using AWS Storage Gateway File Gateway. For the SAP application servers export the VMs as AMIs by using the VM Import/Export feature from AWS. For NFS file shares /sapmnt and /usr/sap/trans establish real-time synchronization from AWS DataSync to Amazon Elastic File System (Amazon EFS).
- C. For the SAP HANA database SAP application servers and NFS file shares use CloudEndure Disaster Recovery to replicate the data continuously from on-premises to AWS. Use CloudEndure Disaster Recovery to launch target instances in the event of a disaster.
- D. For the SAP HANA database use a smaller SAP-certified Amazon EC2 instance. Use SAP HANA system replication with ASYNC replication mode to replicate the data continuously from on-premises to AWS. For the SAP application servers use CloudEndure Disaster Recovery for continuous data replication. For NFS file shares /sapmnt and /usr/sap/trans establish real-time synchronization from AWS DataSync to Amazon Elastic File System (Amazon EFS).

Correct Answer: B

[PAS-C01 PDF Dumps](#)

[PAS-C01 Study Guide](#)

[PAS-C01 Braindumps](#)