



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 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

QUESTION 2

A company's basis administrator is planning to deploy SAP on AWS on Linux. The basis administrator must set up the proper storage to store SAP HANA data and log volumes. Which storage options should the basis administrator choose to meet these requirements? (Select TWO.)

- A. Amazon Elastic Block Store (Amazon EBS) Throughput Optimized HDD (st1)
- B. Amazon Elastic Block Store (Amazon EBS) Provisioned IOPS SSD (io1, k>2)
- C. Amazon S3
- D. Amazon Elastic File System (Amazon EFS)
- E. Amazon Elastic Block Store (Amazon EBS) General Purpose SSD (gp2, gp3)

Correct Answer: CD



QUESTION 3

A company hosts an SAP HANA database on an Amazon EC2 instance in the us-east-1 Region. The company needs to implement a disaster recovery (DR) site in the us-west-1 Region. The company needs a cost-optimized solution that offers a guaranteed capacity reservation an RPO of less than 30 minutes and an RTO of less than 30 minutes.

When solution will meet these requirements?

- A. Deploy a single EC2 instance to support the secondary database in us-west with additional storage Use this secondary database instance to support QA and production Configure the primary SAP HANA database in us-east-1 to constantly replicate the data to the secondary SAP HANA database in us-west-t by using SAP HANA system replication with preload off During DR shut down the QA SAP HANA instance and restart the production services at the secondary site
- B. Deploy a secondary staging server on an EC2 instance in us-west-1 Use CloudEndure Disaster Recovery to replicate changes at the database level from us-east-1 to the secondary staging server on an ongoing basis During DR, initiate cutover increase the size of the secondary EC2 instance to match the primary EC2 instance and start the secondary EC2 instance
- C. Set up the primary SAP HANA database in us-east-1 to constantly replicate the data to a secondary SAP HANA database in us-west-1 by using SAP HANA system replication with preload on Keep the secondary SAP HANA instance as a hot standby that rs ready to take over in case of failure
- D. Create an SAP HANA database AMI by using Amazon Elastic Block Store (Amazon EBS) snapshots Replicate the database and log backup files from a primary Amazon S3 bucket in us-east-1 to a secondary S3 bucket m us-west-1 During DR launch the EC2 instance in us-west-1 based on AMIs that are replicated Update host information Download database and log backups from the secondary S3 bucket Perform a point-in-time recovery

Correct Answer: A

QUESTION 4

A company hosts its SAP NetWeaver workload on SAP HANA m the AWS Cloud The SAP NetWeaver application is protected by a cluster solution that uses Red Hat Enterprise Linux High Availability Add-On The duster solution uses an overlay IP address to ensure that the high availability cluster is still accessible during failover scenarios. An SAP solutions architect needs to facilitate the network connection to this overlay IP address from multiple locations These locations include more than 25 VPCs other AWS Regions and the on-premises environment The company already has set up an AWS Direct Connect connection between the on-premises environment and AWS. What should the SAP solutions architect do to meet these requirements in the MOST scalable manner?

- A. Use VPC peering between the VPCs to route traffic between them
- B. Use AWS Transit Gateway to connect the VPCs and on-premises networks together
- C. Use a Network Load Balancer to route connections to various targets within VPCs
- D. Deploy a Direct Connect gateway to connect the Direct Connect connection over a private VIF to one or more VPCs in any accounts

Correct Answer: D

**QUESTION 5**

A company has deployed SAP HANA in the AWS Cloud. The company needs its SAP HANA database to be highly available. An SAP solutions architect has deployed the SAP HANA database in separate Availability Zones in a single AWS Region. SUSE Linux Enterprise High Availability Extension is configured with an overlay IP address. The overlay IP resource agent has the following IAM policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "oip1",
      "Effect": "Allow",
      "Action": "ec2:AssociateRouteTable",
      "Resource": "arn:aws:ec2:us-east-1:111111111111:route-table/rtb-XYZ"
    },
    {
      "Sid": "oip2",
      "Effect": "Allow",
      "Action": "ec2:DescribeRouteTables",
      "Resource": "*"
    }
  ]
}
```

During a test of failover the SAP solutions architect finds that the overlay IP address does not change to the secondary Availability Zone. Which change should the SAP solutions architect make in the policy statement for Sid oip1 to fix this error?

- A. Change the Action element to ec2 CreateRoute
- B. Change the Action element to ec2 ReplaceRoute
- C. Change the Action element to ec2 ReplaceRouteTableAssociation
- D. Change the Action element to ec2 ReplaceTransitGatewayRoute

Correct Answer: B

[Latest PAS-C01 Dumps](#)

[PAS-C01 PDF Dumps](#)

[PAS-C01 Braindumps](#)