



Free Questions for PAS-C01 by certsinside

Shared by Larsen on 29-01-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

A company plans to migrate its SAP NetWeaver deployment to AWS. The deployment runs on a Microsoft SQL Server database. The company plans to change the source database from SQL Server to SAP HANA as part of this process.

Which migration tools or methods should an SAP solutions architect use to meet these requirements" (Select TWO.)

Options:

- A- SAP HANA classical migration
- B- SAP HANA system replication
- C- SAP Software Update Manager (SUM) Database Migration Option (DMO) with System Move
- D- SAP HANA backup and restore
- E- SAP homogeneous system copy

Answer:

A, C

Explanation:

Question 2

Question Type: MultipleChoice

A company is running an SAP HANA database on AWS. The company wants to manage historical infrequently accessed warm data for a native SAP HANA use case. An SAP solutions architect needs to recommend a solution that can provide online data storage in an extended store available for queries and updates. The solution must be an integrated component of the SAP HANA database and must allow the storage of up to five times more data in the warm tier than in the hot tier.

Which solution will meet these requirements?

Options:

- A-** Use Amazon Data Lifecycle Manager (Amazon DLM) with SAP Data Hub to move data in and out of the SAP HANA database to Amazon S3
- B-** Use an SAP HANA extension node
- C-** Use SAP HANA dynamic tiering as an optional add-on to the SAP HANA database
- D-** Use Amazon Data Lifecycle Manager (Amazon DLM) with SAP HANA spark controller so that SAP HANA can access the data

through the Spark SQL SDA adapter

Answer:

C

Explanation:

This is a feature of the SAP HANA database that allows for the management of historical, infrequently accessed data and provides an integrated solution for data storage in the warm tier. It allows for the storage of up to five times more data in the warm tier than in the hot tier, and it allows for online data storage that is available for queries and updates.

<https://docs.aws.amazon.com/sap/latest/sap-hana/warm-data-tiering-options.html>

Question 3

Question Type: MultipleChoice

A company has an SAP environment that runs on AWS. The company wants to enhance security by restricting Amazon EC2 Instance Metadata Service (IMDS) to IMDSv2 only. The company's current configuration option supports both iMDSv1 and iM0Sv2. The security enhancement must not create an SAP outage.

What should the company do before it applies the security enhancement on EC2 instances that are running the SAP environment?

Options:

- A- Ensure that the SAP kernel versions are 7.45 or later
- B- Ensure that the EC2 instances are Nitro based
- C- Ensure that the AWS Data Provider for SAP is installed on each EC2 instance
- D- Stop the EC2 instances

Answer:

A

Explanation:

A is correct because ensuring that the SAP kernel versions are 7.45 or later is required to support IMDSv2 only. SAP kernel versions below 7.45 do not support IMDSv2 and will fail to start if IMDSv1 is disabled. The other options are not relevant for the security enhancement and may create an SAP outage. Reference: <https://docs.aws.amazon.com/whitepapers/latest/sap-on-aws-technical-deployment-guide/security.html> <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/configuring-instance-metadata-service.html>

Question 4

Question Type: MultipleChoice

A company is planning to migrate its on-premises SAP ERP Central Component (SAP ECC) system on SAP HANA to AWS. Each month the system experiences two peaks in usage. The first peak is on the 21st day of the month when the company runs payroll. The second peak is on the last day of the month when the company processes and exports credit data. Both peak workloads are of high importance and cannot be rescheduled.

The current SAP ECC system has six application servers of a similar size. During normal operation outside of peak usage, four application servers would suffice.

Which purchasing option will meet the company's requirements MOST cost-effectively on AWS?

Options:

- A- Four Reserved Instances and two Spot Instances
- B- Six On-Demand Instances
- C- Six Reserved Instances
- D- Four Reserved Instances and two On-Demand Instances

Answer:

D

Explanation:

D is correct because using four Reserved Instances and two On-Demand Instances provides the most cost-effective purchasing option for the company. Reserved Instances offer lower prices than On-Demand Instances for the four application servers that are needed for normal operation. On-Demand Instances offer flexibility and scalability for the two additional application servers that are needed only during peak usage. Spot Instances are not suitable for high-importance workloads that cannot be rescheduled as they can be interrupted at any time. Reference: <https://docs.aws.amazon.com/whitepapers/latest/sap-on-aws-technical-deployment-guide/cost-optimization.html>
<https://docs.aws.amazon.com/whitepapers/latest/sap-on-aws-technical-deployment-guide/amazon-ec2.html>

Question 5

Question Type: MultipleChoice

A company is hosting an SAP HANA database on AWS. The company is automating operational tasks including backup and system refreshes. The company wants to use SAP HANA Studio to perform data backup of an SAP HANA tenant database to a backint interface. The SAP HANA database is running in multi-tenant database container (MDO mode). The company receives the following error message during an attempt to perform the backup.

```
Could not start backup for system <SID> DBC: [447]: backup could not be completed: [110091] Invalid path selection fo
using backint: /usr/sap/<SID>/SYS/global/hdb/backint/COMPLETE_DATA_BACKUP must start with
/usr/sap/<SID>/SYS/global/hdb/backint/DB_<TENANT>.
```

What should an SAP solutions architect do to resolve this issue?

Options:

- A-** Set the execute permission for AWS Backint agent binary aws-backint-agent and for the launcher script aws-backint-agent-launcher.sh in the installation directory
- B-** Verify the installation steps Create symbolic links (symlinks)
- C-** Ensure that the catalog_backup_using_backint SAP HANA parameter is set to true Ensure that the data_backup_parameter_file and log_backup_parameter_file parameters have the correct path location in the global ini file
- D-** Add the SAP HANA system to SAP HANA Studio Select multiple container mode and then try to initiate the backup again

Answer:

D

Explanation:

https://docs.aws.amazon.com/ja_jp/sap/latest/sap-hana/aws-backint-agent-troubleshooting.html

Question 6

Question Type: MultipleChoice

An SAP solutions architect is leading the SAP basis team for a company. The company's SAP landscape includes SAP HANA database instances for the following systems sandbox development quality assurance test {QAT} system performance test (SPT) and production. The sandbox development and OAT systems are running on Amazon EC2 On-Demand Instances. The SPT and production systems are running on EC2 Reserved Instances All the EC2 instances are using Provisioned IOPS SSO (K)2 Amazon Elastic Block Store (Amazon EBS) volumes.

The entire development learn is in the same time zone and works from 8 AM to 6 PM. The sandbox system is for research and testing that are not critical. The SPT and production systems are business critical The company runs load-testing jobs and stress-testing jobs on the QAT systems overnight to reduce testing duration. The company wants to optimize infrastructure cost for the existing AWS resources.

How can the SAP solutions architect meet these requirements with the LEAST amount of administrative effort?

Options:

A- Use a Spot Fleet instead of the Reserved Instances and On-Demand Instances

B- Use Amazon EventBridge (Amazon CloudWatch Events) and Amazon CloudWatch alarms to stop the development and sandbox EC2 instances from 7 PM every night to 7 AM the next day

C- Make the SAP basis team available 24 hours a day 7 days a week to use the AWS CLI to stop and start the development and sandbox EC2 instances manually

D- Change the EBS volume type to Throughput Optimized HDD (sti) for the /hana/data and nana, log file systems for the production and non-production SAP HANA databases

Answer:

B

Explanation:

B is correct because using Amazon EventBridge (Amazon CloudWatch Events) and Amazon CloudWatch alarms to stop the development and sandbox EC2 instances from 7 PM every night to 7 AM the next day reduces the infrastructure cost for the non-critical systems that are not used during that time. This solution also requires the least amount of administrative effort as it is automated and does not require manual intervention. Reference: <https://docs.aws.amazon.com/whitepapers/latest/sap-on-aws-technical-deployment-guide/cost-optimization.html> <https://docs.aws.amazon.com/whitepapers/latest/sap-on-aws-technical-deployment-guide/amazon-ec2.html>

Question 7

Question Type: MultipleChoice

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?

Options:

- 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 is 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 in 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

Answer:

A

Explanation:

A is correct because deploying a single EC2 instance to support the secondary database in us-west-1 with additional storage is cost-optimized and offers a guaranteed capacity reservation. Configuring the primary SAP HANA database in us-east-1 to constantly replicate the data to the secondary SAP HANA database in us-west-1 by using SAP HANA system replication with preload off provides an RPO of less than 30 minutes. During DR, shutting down the QA SAP HANA instance and restarting the production services at the secondary site provides an RTO of less than 30 minutes. Reference: <https://docs.aws.amazon.com/whitepapers/latest/sap-on-aws-technical-deployment-guide/disaster-recovery.html> <https://docs.aws.amazon.com/whitepapers/latest/sap-on-aws-technical-deployment-guide/sap-hana.html>

Question 8

Question Type: MultipleChoice

A company is planning to migrate its on-premises SAP applications to AWS. The applications are based on Windows operating systems. A file share stores the transport directories and third party application data on the network attached storage of the company's on-premises data center. The company's plan is to lift and shift the SAP applications and the file share to AWS. The company must follow AWS best practices for the migration.

Which AWS service should the company use to host the transport directories and third-party application data on AWS?

Options:

- A- Amazon Elastic Block Store (Amazon EBS)
- B- AWS Storage Gateway
- C- Amazon Elastic File System (Amazon EFS)
- D- Amazon FSx for Windows File Server

Answer:

D

Explanation:

<https://aws.amazon.com/blogs/awsfor-sap/how-to-setup-sap-netweaver-on-windows-mscs-for-sap-ascs-ers-on-aws-using-amazon-fsx/>

Question 9

Question Type: MultipleChoice

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?

Options:

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 .'uv-sap'trans. establish real-time synchronization from AWS DataSync to Amazon Elastic File System (Amazon EFS)

Answer:

D

Explanation:

D is correct because using a smaller SAP certified Amazon EC2 instance for the SAP HANA database reduces the cost of the DR solution. Using SAP HANA system replication with ASYNC replication mode provides an RPO of 15 minutes and an RTO of 4 hours. Using CloudEndure Disaster Recovery for the SAP application servers provides continuous data replication and minimal RTO. Using AWS DataSync to synchronize NFS file shares to Amazon EFS provides real-time data transfer and high availability. Reference:

<https://docs.aws.amazon.com/whitepapers/latest/sap-on-aws-technical-deployment-guide/disaster-recovery.html>

<https://docs.aws.amazon.com/whitepapers/latest/sap-on-aws-technical-deployment-guide/sap-hana.html>

<https://docs.aws.amazon.com/whitepapers/latest/sap-on-aws-technical-deployment-guide/sap-application-servers.html>

<https://docs.aws.amazon.com/whitepapers/latest/sap-on-aws-technical-deployment-guide/nfs-file-shares.html>

Question 10

Question Type: MultipleChoice

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?

Options:

- 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 ReplaceTransrtGatewayRoute

Answer:

B

Explanation:

<https://docs.aws.amazon.com/sap/latest/sap-hana/sap-hana-on-aws-cluster-configuration-prerequisites.html>

To Get Premium Files for PAS-C01 Visit

<https://www.p2pexams.com/products/pas-c01>

For More Free Questions Visit

<https://www.p2pexams.com/amazon/pdf/pas-c01>

