



Vendor: Amazon

Exam Code: PAS-C01

Exam Name: AWS Certified: SAP on AWS - Specialty

Version: DEMO

QUESTION 1

A company has deployed SAP workloads on AWS. The AWS Data Provider for SAP is installed on the Amazon EC2 instance where the SAP application is running. An SAP solutions architect has attached an IAM role to the EC2 instance with the following policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AWSDataProvider1",
      "Effect": "Allow",
      "Action": [
        "EC2:DescribeInstances",
        "EC2:DescribeVolumes"
      ],
      "Resource": "*"
    },
    {
      "Sid": "AWSDataProvider2",
      "Effect": "Allow",
      "Action": "s3:GetObject",
      "Resource": [
        "arn:aws:s3:::aws-sap-data-provider/config.properties"
      ]
    }
  ]
}
```

The AWS Data Provider for SAP is not returning any metrics to the SAP application. Which change should the SAP solutions architect make to the IAM permissions to resolve this issue?

- A. Add the cloudwatch:ListMetrics action to the policy statement with Sid AWSDataProvider1.
- B. Add the cloudwatch:GetMetricStatistics action to the policy statement with Sid AWSDataProvider1.
- C. Add the cloudwatch:GetMetricStream action to the policy statement with Sid AWSDataProvider1.
- D. Add the cloudwatch:DescribeAlarmsForMetric action to the policy statement with Sid AWSDataProvider1.

Answer: B

Explanation:

<https://docs.aws.amazon.com/sap/latest/general/data-provider-troubleshooting.html>

QUESTION 2

A company wants to deploy an SAP HANA database on AWS by using AWS Launch Wizard for SAP. An SAP solutions architect needs to run a custom post-deployment script on the Amazon EC2 instance that Launch Wizard provisions.

Which actions can the SAP solutions architect take to provide the post-deployment script in the Launch Wizard console? (Choose two.)

- A. Provide the FTP URL of the script.
- B. Provide the HTTPS URL of the script on a web server.

- C. Provide the Amazon S3 URL of the script.
- D. Write the script inline.
- E. Upload the script.

Answer: CE

Explanation:

<https://catalog.us-east-1.prod.workshops.aws/workshops/754ba343-2704-404a-8abe-be7b21c4d9d5/en-US/800-other/802-prepostscript>

QUESTION 3

A company is planning to move its on-premises SAP HANA database to AWS. The company needs to migrate this environment to AWS as quickly as possible. An SAP solutions architect will use AWS Launch Wizard for SAP to deploy this SAP HANA workload.

Which combination of steps should the SAP solutions architect follow to start the deployment of this workload on AWS? (Choose three.)

- A. Download the SAP HANA software.
- B. Download the AWS CloudFormation template for the SAP HANA deployment.
- C. Download and extract the SAP HANA software. Upload the SAP HANA software to an FTP server that Launch Wizard can access.
- D. Upload the unextracted SAP HANA software to an Amazon S3 destination bucket. Follow the S3 file path syntax for the software in accordance with Launch Wizard recommendations.
- E. Bring the operating system AMI by using the Bring Your Own Image (BYOI) model, or purchase the subscription for the operating system AMI from AWS Marketplace.
- F. Create the SAP file system by using Amazon Elastic Block Store (Amazon EBS) before the deployment.

Answer: ADE

Explanation:

<https://docs.aws.amazon.com/launchwizard/latest/userguide/launch-wizard-sap-setting-up.html>
<https://docs.aws.amazon.com/launchwizard/latest/userguide/launch-wizard-sap-structure.html>

QUESTION 4

A company wants to implement SAP HANA on AWS with the Multi-AZ deployment option by using AWS Launch Wizard for SAP. The solution will use SUSE Linux Enterprise High Availability Extension for the high availability deployment. An SAP solutions architect must ensure that all the prerequisites are met. The SAP solutions architect also must ensure that the user inputs to start the guided deployment of Launch Wizard are valid.

Which combination of steps should the SAP solutions architect take to meet these requirements? (Choose two.)

- A. Before starting the Launch Wizard deployment, create the underlying Amazon Elastic Block Store (Amazon EBS) volume types to use for SAP HANA data and log volumes based on the performance requirements.
- B. Use a value for the PaceMakerTag parameter that is not used by any other Amazon EC2 instances in the AWS Region where the system is being deployed.
- C. Ensure that the virtual hostname for the SAP HANA database that is used for the SUSE Linux Enterprise High Availability Extension configuration is not used in any other deployed accounts.
- D. Ensure that the VirtualIPAddress parameter is outside the VPC CIDR and is not being used in the route table that is associated with the subnets where primary and secondary SAP HANA instances will be deployed.
- E. Before starting the Launch Wizard deployment, set up the SUSE Linux Enterprise High

Availability Extension network configuration and security group.

Answer: BD

Explanation:

<https://docs.aws.amazon.com/launchwizard/latest/userguide/launch-wizard-sap-deploying.html#launch-wizard-hana>

QUESTION 5

A company is running its SAP workload on AWS. The company's security team has implemented the following requirements:

- All Amazon EC2 instances for SAP must be SAP certified instance types.
- Encryption must be enabled for all Amazon S3 buckets and Amazon Elastic Block Store (Amazon EBS) volumes.
- AWS CloudTrail must be activated.
- SAP system parameters must be compliant with business rules.
- Detailed monitoring must be enabled for all instances.

The company wants to develop an automated process to review the systems for compliance with the security team's requirements. The process also must provide notification about any deviation from these standards.

Which solution will meet these requirements?

- A. Use AWS AppConfig to model configuration data in an AWS Systems Manager Automation runbook. Schedule this Systems Manager Automation runbook to monitor for compliance with all the requirements. Integrate AWS AppConfig with Amazon CloudWatch for notification purposes.
- B. Use AWS Config managed rules to monitor for compliance with all the requirements. Use Amazon EventBridge (Amazon CloudWatch Events) and Amazon Simple Notification Service (Amazon SNS) for email notification when a resource is flagged as noncompliant.
- C. Use AWS Trusted Advisor to monitor for compliance with all the requirements. Use Trusted Advisor preferences for email notification when a resource is flagged as noncompliant.
- D. Use AWS Config managed rules to monitor for compliance with the requirements, except for the SAP system parameters. Create AWS Config custom rules to validate the SAP system parameters. Use Amazon EventBridge (Amazon CloudWatch Events) and Amazon Simple Notification Service (Amazon SNS) for email notification when a resource is flagged as noncompliant.

Answer: D

Explanation:

<https://aws.amazon.com/blogs/awsfor sap/audit-your-sap-systems-with-aws-config-part-ii/>

QUESTION 6

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.

Answer: B

Explanation:

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

QUESTION 7

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?

- 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 8

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.

Which solution will meet these requirements?

- A. Deploy a single EC2 instance to support the secondary database in us-west-1 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-1 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: B

Explanation:

CloudEndure Disaster Recovery continuously replicates your machines (including operating system, system state configuration, databases, applications, and files) into a low-cost staging area in your target AWS account and preferred Region.

<https://aws.amazon.com/blogs/awsfor/sap/sap-disaster-recovery-solution-using-cloudendure-part-1-failover/>

QUESTION 9

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 (MDC) 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 for data backup 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?

- 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 10

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, all 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?

- 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:

Four Reserved Instances and two on demand Instances would meet the company's requirements most cost-effectively on AWS. The company can use four Reserved Instances for the normal operation and two on demand Instances for the peak workloads during the 21st day of the month and the last day of the month. This will provide the required capacity at the lowest cost.

QUESTION 11

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 IMDSv2. 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?

- 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: B

Explanation:

<https://me.sap.com/notes/1656250>

QUESTION 12

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 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?

- 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:

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

QUESTION 13

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? (Choose two.)

- 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: AC

Explanation:

<https://docs.aws.amazon.com/sap/latest/sap-hana/migrating-hana-anydb-to-hana.html>

QUESTION 14

An SAP basis architect is configuring high availability for a critical SAP system on AWS. The SAP basis architect is using an overlay IP address to route traffic to the subnets across multiple Availability Zones within an AWS Region for the system's SAP HANA database.

What should the SAP basis architect do to route the traffic to the Amazon EC2 instance of the active SAP HANA database?

- A. Edit the route in the route table of the VPC that includes the EC2 instance that runs SAP HANA. Specify the overlay IP address as the destination. Specify the private IP address of the EC2 instance as the target.
- B. Edit the inbound and outbound rules in the security group of the EC2 instance that runs SAP HANA. Allow traffic for SAP HANA specific ports from the overlay IP address.
- C. Edit the network ACL of the subnet that includes the EC2 instance that runs SAP HANA. Allow traffic for SAP HANA specific ports from the overlay IP address.
- D. Edit the route in the route table of the VPC that includes the EC2 instance that runs SAP HANA. Specify the overlay IP address as the destination. Specify the elastic network interface of the EC2 instance as the target.

Answer: D

Explanation:

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

prerequisites.html

QUESTION 15

A company is running SAP ERP Central Component (SAP ECC) with a Microsoft SQL Server database on AWS. A solutions architect must attach an additional 1 TB Amazon Elastic Block Store (Amazon EBS) volume. The company needs to write the SQL Server database backups to this EBS volume before moving the database backups to Amazon S3 for long-term storage. Which EBS volume type will meet these requirements MOST cost-effectively?

- A. Throughput Optimized HDD (st1)
- B. Provisioned IOPS SSD (io2)
- C. General Purpose SSD (gp3)
- D. Cold HDD (sc1)

Answer: D

Explanation:

<https://docs.aws.amazon.com/sap/latest/sap-hana/cold-data-tiering-options.html#sap-archiving>
For archiving, another option is to use the Amazon Elastic Block Store (Amazon EBS) sc1 volume type as the underlying storage type for your archive file system. Amazon EBS sc1 volumes are inexpensive block storage and are designed for less frequently accessed workloads like data archiving. To increase durability and availability of your archived data, we recommend that you copy the data to Amazon S3 for backup and Amazon S3 Glacier for long term retention.

QUESTION 16

A company is moving to the AWS Cloud gradually. The company has multiple SAP landscapes on VMware. The company already has sandbox, development, and QA systems on AWS. The company's production system is still running on premises. The company has 2 months to cut over the entire landscape to the AWS Cloud.

The company has adopted a hybrid architecture for the next 2 months and needs to synchronize its shared file systems between the landscapes. These shared file systems include /trans directory mounts, /software directory mounts, and third-party integration mounts. In the on-premises landscape, the company has NFS mounts between the servers. On the AWS infrastructure side, the company is using Amazon Elastic File System (Amazon EFS) to share the common files.

An SAP solutions architect needs to design a solution to schedule transfer of these shared files bidirectionally four times each day. The data transfer must be encrypted.

Which solution will meet these requirements?

- A. Write an rsync script. Schedule the script through cron for four times each day in the on-premises VMware servers to transfer the data from on premises to AWS.
- B. Install an AWS DataSync agent on the on-premises VMware platform. Use the DataSync endpoint to synchronize between the on-premises NFS server and Amazon EFS on AWS.
- C. Order an AWS Snowcone device. Use the Snowcone device to transfer data between the on-premises servers and AWS.
- D. Set up a separate AWS Direct Connect connection for synchronization between the on-premises servers and AWS.

Answer: B

Explanation:

<https://aws.amazon.com/datasync/faqs/>

Thank You for Trying Our Product

Passleader Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.passleader.com/all-products.html>



Microsoft



ORACLE



JUNIPER
NETWORKS



EMC²
where information lives

10% Discount Coupon Code: ASTR14