

Exam	PAS-C01
Title	Amazon AWS Certified: SAP on AWS - Specialty Exam
Version	2.0
Product Type	65 Q&A with explanations

QUESTION 1

A global enterprise is running SAP ERP Central Component (SAP ECC) workloads on Oracle in an on-premises environment. The enterprise plans to migrate to SAP S/4HANA on AWS.

The enterprise recently acquired two other companies. One of the acquired companies is running SAP ECC on Oracle as its ERP system. The other acquired company is running an ERP system that is not from SAP. The enterprise wants to consolidate the three ERP systems into one ERP system on SAP S/4HANA on AWS. Not all the data from the acquired companies needs to be migrated to the final ERP

system. The enterprise needs to complete this migration with a solution that minimizes cost and maximizes operational efficiency.

Which solution will meet these requirements?

- A. Perform a lift-and-shift migration of all the systems to AWS. Migrate the ERP system that is not from SAP to SAP ECC. Convert all three systems to SAP S/4HANA by using SAP Software Update Manager (SUM) Database Migration Option (DMO). Consolidate all three SAP S/4HANA systems into a final SAP S/4HANA system. Decommission the other systems.
- B. Perform a lift-and-shift migration of all the systems to AWS. Migrate the enterprise's initial system to SAP HANA, and then perform a conversion to SAP S/4HANA. Consolidate the two systems from the acquired companies with this SAP S/4HANA system by using the Selective Data Transition approach with SAP Data Management and Landscape Transformation (DMLT).
- C. Use SAP Software Update Manager (SUM) Database Migration Option (DMO) with System Move to re-architect the enterprise initial system to SAP S/4HANA and to change the platform to AWS. Consolidate the two systems from the acquired companies with this SAP S/4HANA system by using the Selective Data Transition approach with SAP Data Management and Landscape Transformation (DMLT).
- D. Use SAP Software Update Manager (SUM) Database Migration Option (DMO) with System Move to re-architect all the systems to SAP S/4HANA and to change the platform to AWS. Consolidate all three SAP S/4HANA systems into a final SAP S/4HANA system. Decommission the other systems.

Answer: A

QUESTION 2

A global retail company is running its SAP landscape on AWS. Recently, the company made changes to its SAP Web Dispatcher architecture. The company added an additional SAP Web Dispatcher for high availability with an Application Load Balancer (ALB) to balance the load between the two SAP Web Dispatchers. When users try to access SAP through the ALB, the system is reachable. However, the SAP backend system is showing an error message. An investigation reveals that the issue is related to SAP session handling and distribution of requests. The company confirmed that the system was working as

expected with one SAP Web Dispatcher. The company replicated the configuration of that SAP Web Dispatcher to the new SAP Web Dispatcher.

How can the company resolve the error?

- A. Maintain persistence by using session cookies. Enable session stickiness (session affinity) on the SAP Web Dispatchers by setting the `wdisp/HTTP/esid_support` parameter to True.
- B. Maintain persistence by using session cookies. Enable session stickiness (session affinity) on the ALB.
- C. Turn on host-based routing on the ALB to route traffic between the SAP Web Dispatchers.
- D. Turn on URL-based routing on the ALB to route traffic to the application based on URL.

Answer: C

QUESTION 3

A company hosts its SAP NetWeaver workload on SAP HANA in the AWS Cloud. The SAP NetWeaver application is protected by a cluster solution that uses Red Hat Enterprise Linux High Availability Add-On. The cluster 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 in 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

Answer: D

QUESTION 4

A company is implementing SAP HANA on AWS According to the company's security policy SAP backups must be encrypted Only authorized team members can have the ability to decrypt the SAP backups What is the MOST operationally efficient solution that meets these requirements?

- A. Configure AWS Backup Agent for SAP HANA to create SAP backups in an Amazon S3 bucket After a backup is created encrypt the backup by using client-side encryption Share the encryption key with authorized team members only
- B. Configure AWS Backup Agent for SAP HANA to use AWS Key Management Service (AWS KMS) for SAP backups Create a key policy to grant decryption permission to authorized team members only
- C. Configure AWS Storage Gateway to transfer SAP backups from a file system to an Amazon S3 bucket Use an S3 bucket policy to grant decryption permission to authorized team members only
- D. Configure AWS Backup Agent for SAP HANA to use AWS Key Management Service (AWS KMS) for SAP backups Grant object ACL decryption permission to authorized team members only

Answer: C

QUESTION 5

A data analysis company has two SAP landscapes that consist of sandbox development QA, preproduction and production servers. One landscape is on Windows and the other landscape is on Red Hat Enterprise Linux. The servers reside in a room in a building that other tenants share.

An SAP solutions architect proposes to migrate the SAP applications to AWS The SAP solutions architect wants to move the production backups to AWS and wants to make the backups highly available to restore in case of unavailability of an on-premises server.

Which solution will meet these requirements MOST cost-effectively?

- A. Take a backup of the production servers Implement an AWS Storage Gateway Volume Gateway Create file shares by using the Storage Gateway Volume Gateway Copy the backup files to the file shares through NFS and SMB
- B. Take a backup of the production servers Send those backups to tape drives implement an AWS Storage Gateway Tape Gateway Send the backups to Amazon S3 Standard-Infrequent Access (S3 Standard-IA) through the S3 console Move the backups immediately to S3 Glacier Deep Archive
- C. Implement a third-party tool to take images of the SAP application servers and database server Take regular snapshots at 1-hour intervals send the snapshots to Amazon S3 Glacier directly through the S3 Glacier console Store the same images in different S3 buckets in different AWS Regions
- D. Take a backup of the production servers Implement an Amazon S3 File Gateway Create file shares by using the S3 File Gateway Copy the backup files to the file shares through NFS and SMB Map backup files directly to Amazon S3 Configure an S3 Lifecycle policy to send the backup files to S3 Glacier based on the company's data retention policy

Answer: C

QUESTION 6

A company's SAP basis team is responsible for database backups in Amazon S3. The company frequently needs to restore the last 3 months of backups into the pre-production SAP system to perform tests and analyze performance. Previously an employee accidentally deleted backup files from the S3 bucket. The SAP basis team wants to prevent accidental deletion of backup files in the future.

Which solution will meet these requirements?

- A. Create a new resource-based policy that prevents deletion of the S3 bucket

- B. Enable versioning and multi-factor authentication (MFA) on the S3 bucket
- C. Create signed cookies for the backup files in the S3 bucket Provide the signed cookies to authorized users only
- D. Apply an S3 Lifecycle policy to move the backup files immediately to S3 Glacier

Answer: A

QUESTION 7

A company wants to run SAP HANA on AWS in the eu-central-1 Region. The company must make the SAP HANA system highly available by using SAP HANA system replication in addition the company must create a disaster recovery (DR) solution that uses SAP HANA system replication in the eu-west-1 Region As prerequisites the company has confirmed that inter-AZ latency is less than 1 ms and that Inter-Region latency is greater than 1 ms.

Which solutions will meet these requirements? (Select TWO.)

- A. Install the tier 1 primary system and the tier 2 secondary system in eu-central-1 Configure the tier 1 system in Availability Zone 1 Configure the tier 2 system in Availability Zone 2 Configure SAP HANA system replication between tier 1 and tier 2 by using ASYNC replication mode install the OR tier 3 secondary system in eu-west-1 by using SYNC replication mode.
- B. Install the tier 1 primary system and the tier 2 secondary system in eu-central-1 Configure the tier 1 system in Availability Zone 1 Configure the tier 2 system in Availability Zone 2 Configure SAP HANA system replication between tier 1 and tier 2 by using SYNC replication mode Install the OR tier 3 secondary system in eu-west-1 by using ASYNC replication mode.
- C. Install the tier 1 primary system and the tier 2 secondary system in eu-central-1 Configure the tier 1 system in Availability Zone 1 Configure the tier 2 system in Availability Zone 2 Configure SAP HANA system replication between tier 1 and tier 2 by using SYNC replication mode Install the OR tier 3 secondary system in eu-west-1 Store daily backups from tier 1 in an Amazon S3 bucket in eu-central-1 Use S3 Cross-Region Replication to copy the daily backups to eu-west-1 where they can be restored if needed
- D. Install the tier 1 primary system in eu-central-1 install the tier 2 secondary system and the DR tier 3 secondary system in eu-west-1 Configure the tier 2 system in Availability Zone 1 Configure the tier 3 system in Availability Zone 2 Configure SAP HANA system replication between all tiers by using ASYNC replication mode
- E. Install the tier 1 primary system and the tier 2 secondary system in eu-central-1 Configure the tier 1 system in Availability Zone 1 Configure the tier 2 system in Availability Zone 2 Configure SAP HANA system replication between tier 1 and tier 2 by using SYNCMEM replication mode Install the DR tier 3 secondary system in eu-west-1 by using ASYNC replication mode

Answer: BC

QUESTION 8

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 Backint Agent for SAP HANA (AWS Backint 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 AWS Backint 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 AWS Backint agent increase this value in the awsbackint-agent-config.yaml configuration file based on the EC2 instance type and storage configurations
- C. Check the MaximumConcurrentFilesForRestore parameter for AWS Backint agent Increase the parameter from 5 to 10 by using the aws-backint-agent-config.yaml configuration file
- D. Ensure that the backups are compressed if necessary configure AWS Backint agent to compress the backups and send them to an Amazon S3 bucket

Answer: A

QUESTION 9

A company wants to migrate its SAP workloads to AWS from another cloud provider. The company's landscape consists of SAP S/4HANA SAP BW/4HANA SAP Solution Manager and SAP Web Dispatcher SAP Solution Manager 15 running on

SAP NANA

The company wants to change the operating system from SUSE Linux Enterprise Server to Red Hat Enterprise Linux as a part of this migration. The company needs a solution that results in the least possible downtime for the SAP S/4HANA and SAP BW 4HANA systems.

Which migration solution will meet these requirements?

- A. Use SAP Software Provisioning Manager to perform a system export/import for SAP S/4HANA SAP BW 4HANA SAP Solution Manager and SAP Web Dispatcher
- B. Use backup and restore for SAP S/4HANA SAP BW 4HANA and SAP Solution Manager Reinstall SAP Web Dispatcher on AWS with the necessary configuration
- C. Use backup and restore for SAP S/4HANA and SAP BW 4HANA Use SAP Software Provisioning Manager to perform a system export/import for SAP Solution Manager Reinstall SAP Web Dispatcher on AWS with the necessary configuration.
- D. Use SAP HANA system replication to replicate the data between the source system and the target AWS system for SAP S/4HANA and SAP BW 4HANA Use SAP Software Provisioning Manager to perform a system export/import for SAP Solution Manager Reinstall SAP Web Dispatcher on AWS with the necessary configuration

Answer: D

QUESTION 10

A company is running an SAP on Oracle system on IBM Power architecture in an on-premises data center. The company wants to migrate the SAP system to AWS. The Oracle database is 15 TB in size. The company has set up a 100 Gbps AWS Direct Connect connection to AWS from the on-premises data center. Which solution should the company use to migrate the SAP system MOST quickly?

- A. Before the migration window build a new installation of the SAP system on AWS by using SAP Software Provisioning Manager. During the migration window export a copy of the SAP system and database by using the heterogeneous system copy process and R3load Copy the output of the SAP system files to AWS through the Direct Connect connection import the SAP system to the new SAP installation on AWS Switch over to the SAP system on AWS
- B. Before the migration window build a new installation of the SAP system on AWS by using SAP Software Provisioning Manager Back up the Oracle database by using native Oracle tools Copy the backup of the Oracle database to AWS through the Direct Connect connection Import the Oracle database to the SAP system on AWS Configure Oracle Data Guard to begin replication on-premises database log changes from the SAP system to the new AWS system During the migration window use Oracle to replicate any remaining changes to the Oracle database hosted on AWS Switch over to the SAP system on AWS
- C. Before the migration window build a new installation of the SAP system on AWS by using SAP Software Provisioning Manager Create a staging Oracle database on premises to perform Cross Platform Transportable Tablespace (XTTS) conversion on the Oracle database Take a backup of the converted staging database Copy the converted backup to AWS through the Direct Connect connection import the Oracle database backup to the SAP system on AWS Take regularly scheduled incremental backups and XTTS conversions of the staging database Transfer these backups and conversions to the AWS target database During the migration window, perform a final incremental Oracle backup Convert the final Oracle backup by using XTTS Replay the logs in the target Oracle database hosted on AWS Switch over to the SAP system on AWS
- D. Before the migration window launch an appropriately sized Amazon EC2 instance on AWS to receive the migrated SAP database Create an AWS Server Migration Service (AWS SMS) job to take regular snapshots of the on-premises Oracle hosts Use AWS SMS to copy the snapshot as an AMI to AWS through the Direct Connect connection Create a new SAP on Oracle system by using the migrated AMI During the migration window take a final incremental SMS snapshot and copy the snapshot to AWS Restart the SAP system by using the new up-to-date AMI Switch over to the SAP system on AWS

Answer: D

QUESTION 11

An SAP solutions architect is designing an SAP HANA scale-out architecture for SAP Business Warehouse (SAP BW) on SAP HANA on AWS. The SAP solutions architect identifies the design as a three-node scale out deployment of x1e.32xlarge Amazon EC2 instances

certexam.org

The SAP solutions architect must ensure that the SAP HANA scale-out nodes can achieve the lowlatency and high-

throughput network performance that are necessary for node-to-node communication

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

- A. Create a cluster placement group Launch the instances into the cluster placement group
- B. Create a spread placement group Launch the instances into the spread placement group
- C. Create a partition placement group Launch the instances into the partition placement group
- D. Based on the operating system version verify that enhanced networking is enabled on all the nodes
- E. Switch to a different instance family that provides network throughput that is greater than 25 Gbps

Answer: CE

QUESTION 12

A company needs to migrate its critical SAP workloads from an on-premises data center to AWS. The company has a few source production databases that are 10 TB or more in size. The company wants to minimize the downtime for this migration. As part of the proof of concept, the company used a low-speed, high-latency connection between its data center and AWS. During the actual migration, the company wants to maintain a consistent connection that delivers high bandwidth and low latency. The company also wants to add a layer of connectivity resiliency. The backup connectivity does not need to be as fast as the primary connectivity.

An SAP solutions architect needs to determine the optimal network configuration for data transfer.

The solution must transfer the data with minimum latency.

Which configuration will meet these requirements?

- A. Set up one AWS Direct Connect connection for connectivity between the on-premises data center and AWS. Add an AWS Site-to-Site VPN connection as a backup to the Direct Connect connection.
- B. Set up an AWS Direct Connect gateway with multiple Direct Connect connections that use a link aggregation group (LAG) between the on-premises data center and AWS.
- C. Set up Amazon Elastic File System (Amazon EFS) file system storage between the on-premises data center and AWS. Configure a cron job to copy the data into this EFS mount. Access the data in the EFS file system from the target environment.
- D. Set up two redundant AWS Site-to-Site VPN connections for connectivity between the on-premises data center and AWS.

Answer: A

QUESTION 13

A company wants to migrate its SAP ERP landscape to AWS. The company will use a highly available, distributed deployment for the new architecture. Clients will access SAP systems from a local data center through an AWS Site-to-Site VPN connection that is already in place. An SAP solutions architect needs to design the network access to the SAP production environment.

Which configuration approaches will meet these requirements? (Select TWO.)

- A. For the ASCS instance, configure an overlay IP address that is within the production VPC CIDR range. Create an AWS Transit Gateway. Attach the VPN to the transit gateway. Use the transit gateway to route the communications between the local data center and the production VPC. Create a static route on the production VPC to route traffic that is directed to the overlay IP address to the ASCS instance.
- B. For the ASCS instance, configure an overlay IP address that is outside the production VPC CIDR range. Create an AWS Transit Gateway. Attach the VPN to the transit gateway. Use the transit gateway to route the communications between the local data center and the production VPC. Create a static route on the production VPC to route traffic that is directed to the overlay IP address to the ASCS instance.
- C. For the ASCS instance, configure an overlay IP address that is within the production VPC CIDR range. Create a target group that points to the overlay IP address. Create a Network Load Balancer and register the target group. Create a static route on the production VPC to route traffic that is directed to the overlay IP address to the ASCS instance.
- D. For the ASCS instance, configure an overlay IP address that is outside the production VPC CIDR range. Create a target group that points to the overlay IP address. Create a Network Load Balancer, and register the target group. Create a static route on the production VPC to route traffic that is directed to the overlay IP address to the ASCS instance.
- E. For the ASCS instance, configure an overlay IP address that is outside the production VPC CIDR range. Create a target group that points to the overlay IP address. Create an Application Load Balancer and register the target group. Create a static