Amazon

(AWS Certified SAP on AWS - Specialty PAS-C01)

AWS Certified SAP on AWS - Specialty PAS-C01

Total: **130 Questions**

Link:

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’s 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: C

Explanation:

The correct answer is C. Here's why:

Why C is the best solution:

Minimizes cost and maximizes operational efficiency: Option C leverages the Selective Data Transition approach for the acquired companies. This is crucial because the requirement states that "Not all the data from the acquired companies needs to be migrated." A Selective Data Transition approach only migrates the necessary data, leading to faster migrations, lower storage costs, and a cleaner SAP S/4HANA system. SAP DMLT supports this approach.
Direct S/4HANA conversion: Option C converts the enterprise's initial system directly to SAP S/4HANA on AWS using DMO with System Move. This avoids an intermediate step like migrating to SAP ECC as proposed in option A or migrating to SAP HANA first and then converting to S/4HANA as suggested in option B. Fewer steps reduce migration time and complexity, thus reducing cost and increasing operational efficiency. DMO with System Move is designed for such platform migrations.
Avoids unnecessary full system conversions: Option D suggests converting all systems using DMO. However, given the requirement to migrate only select data from acquired companies, converting entire systems only to consolidate them later is inefficient and costly.
Leverages SAP best practices: DMO with System Move and Selective Data Transition (using tools like SAP DMLT) are recognized SAP methodologies for system conversions and data migration.

Why other options are less suitable:

Option A: Performing a lift-and-shift of all systems and then migrating the non-SAP system to SAP ECC adds unnecessary complexity and cost. Migrating everything and then consolidating is inefficient. The requirement specified there is a non-SAP system that will be migrating to S/4HANA but did not mention that it needed to be migrated to ECC first.
Option B: First migrating the main system to SAP HANA and then converting it to S/4HANA is a viable

approach but less efficient than directly converting to S/4HANA as offered in option C. Moreover, a lift-and-shift approach to all systems is unnecessary and more costly than needed.

Option D: Converting all systems to SAP S/4HANA only to consolidate them later is an inefficient and wasteful approach, especially since not all data needs to be migrated.

Supporting Resources:

SAP S/4HANA System Conversion:https://www.sap.com/services/s4hana-migration.html
SAP Data Management and Landscape Transformation (DMLT):https://www.sap.com/services/data-landscape-transformation.html
SAP Software Update Manager (SUM):https://support.sap.com/en/tools/software-lifecycle-management/software-update-manager.html
AWS for SAP:https://aws.amazon.com/sap/

Therefore, option C provides the most cost-effective and operationally efficient solution by directly converting the primary system to S/4HANA and selectively migrating data from the acquired companies, aligning with the requirements of the scenario.

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: B

Explanation:

The correct solution is **B. Maintain persistence by using session cookies. Enable session stickiness (session affinity) on the ALB.**

Here's why:

The problem describes a scenario where adding a second SAP Web Dispatcher behind an Application Load Balancer (ALB) is causing session handling issues in the SAP backend. This indicates that subsequent requests from the same user are not consistently being routed to the same SAP Web Dispatcher, leading to loss of session context.

SAP relies on maintaining session state for each user interaction. Without session stickiness, the ALB might route different requests from the same user to different SAP Web Dispatchers. Each Web Dispatcher would then treat the request as a new session, leading to errors because the SAP backend expects continuity within a session.

Enabling session stickiness (also known as session affinity) on the ALB solves this problem. When stickiness is

enabled, the ALB uses cookies to track which target (in this case, SAP Web Dispatcher) a client is associated with. Subsequent requests from the same client (identified by the cookie) are then consistently routed to the same target.

Option A is incorrect because while wdisp/HTTP/esid_support on SAP Web Dispatcher is important for session ID management, it doesn't guarantee stickiness at the ALB level. The ALB still needs to be configured to respect and enforce that affinity.

Option C (host-based routing) and D (URL-based routing) are not suitable for this scenario. These routing methods are useful for directing traffic based on the hostname or URL path, but they don't maintain session affinity. The goal here is to ensure that all requests from a specific user session consistently reach the same SAP Web Dispatcher.

In summary, the ALB needs to maintain stickiness to ensure that all requests from the same user session are directed to the same SAP Web Dispatcher. Configuring session stickiness on the ALB, using session cookies to maintain persistence, achieves this effectively.

Relevant Documentation:

Application Load Balancers: Sticky Sessions:
https://docs.aws.amazon.com/elasticloadbalancing/latest/application/sticky-sessions.html
SAP Web Dispatcher:
https://help.sap.com/docs/SAP_NETWEAVER_AS_ABAP_752/4fc4a8818a9e456bb29a099d00f1b64f/476889ae6f

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, 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: B

Explanation:

The requirement is to provide scalable network connectivity to an overlay IP address used by a high availability SAP cluster across multiple VPCs, AWS Regions, and an on-premises environment.

Option B, using AWS Transit Gateway, is the most scalable and efficient solution. Transit Gateway acts as a central hub, simplifying network management and reducing the complexity of connecting numerous VPCs and on-premises networks. It enables transitive routing, allowing traffic to flow between connected networks. The scalability of Transit Gateway makes it suitable for environments with a large number of VPCs and diverse connection types (VPCs, Direct Connect).

Option A, using VPC peering, becomes complex and unmanageable with more than a few VPCs due to the need for individual peering connections between each VPC, leading to a full mesh configuration, which is difficult to maintain and scale. VPC peering also doesn't inherently address connectivity to other AWS

Regions or the on-premises environment via Direct Connect.

Option C, using a Network Load Balancer (NLB), is designed for distributing traffic to multiple targets within VPCs. While NLBs can handle high availability within a VPC, they don't inherently solve the cross-VPC, cross-Region, and on-premises connectivity challenges described in the scenario. The NLB's primary purpose is load balancing, not routing between networks.

Option D, deploying a Direct Connect gateway, primarily focuses on connecting Direct Connect to VPCs. While necessary for on-premises connectivity, it doesn't address the connectivity between the multiple VPCs within AWS or connections to other AWS Regions. Direct Connect gateway works in conjunction with a solution like Transit Gateway, not as a replacement for it in this scenario.

In summary, AWS Transit Gateway's hub-and-spoke architecture provides the most scalable and manageable solution for connecting multiple VPCs, AWS Regions, and on-premises environments via Direct Connect, making it the best choice to facilitate network connections to the overlay IP address of the SAP high availability cluster.

AWS Transit Gateway DocumentationAWS Direct Connect Gateway Documentation

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 Backint 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 Backint 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 Backint 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: B

Explanation:

Option B is the most operationally efficient and secure solution for encrypting SAP HANA backups in AWS while adhering to the company's security policy of restricted decryption access.

Here's why:

AWS Key Management Service (KMS) Integration: KMS is a managed service specifically designed for creating, managing, and controlling cryptographic keys. Integrating Backint Agent with KMS allows for seamless encryption and decryption of backups directly within the AWS ecosystem.

Centralized Key Management: KMS provides a centralized and auditable location for managing encryption keys. This simplifies key rotation, access control, and compliance.

Granular Access Control: KMS key policies enable fine-grained control over who can decrypt backups. By granting decryption permissions only to authorized team members, the company enforces its security policy effectively.

Operational Efficiency: KMS integration eliminates the need for manual key management and distribution, which can be complex and error-prone. Backups are automatically encrypted during the backup process, and decryption is handled securely by AWS when authorized users access the backup.

Option A requires client-side encryption, which would involve additional overhead in key management and distribution. Sharing the encryption key outside of AWS managed services introduces risk and is not as operationally efficient as using KMS.

Option C involves AWS Storage Gateway, which is unnecessary and adds complexity. Backint Agent can directly back up to S3 without the need for a gateway. Furthermore, S3 bucket policies are not the ideal mechanism for managing decryption permissions for encryption performed by Backint; KMS policies offer more granular control.

Option D uses object ACLs. While object ACLs can control access to S3 objects, they are not best practice for decryption of KMS-encrypted objects. KMS key policies are the preferred and more granular way to control access for KMS-based encryption.

In summary, Option B leverages the strengths of AWS KMS for centralized key management, granular access control, and operational efficiency, making it the most appropriate solution.

Supporting links:

AWS KMS: https://aws.amazon.com/kms/
AWS Backint Agent for SAP HANA: https://aws.amazon.com/sap/partners/backint/
KMS Key Policies: https://docs.aws.amazon.com/kms/latest/developerguide/key-policies.html

Question: 5

A data analysis company has two SAP landscapes that consist of sandbox, development, QA, pre-production, 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: D

Explanation:

Here's a detailed justification for why option D is the most cost-effective solution for backing up SAP production data to AWS, ensuring high availability and adhering to data retention policies:

Option D leverages Amazon S3 File Gateway, a hybrid cloud storage service, which presents a local, on-premises interface (NFS/SMB file shares) to the SAP systems. This allows seamless copying of backup files to the gateway. The gateway then asynchronously and efficiently transfers the data to Amazon S3. A key cost-saving aspect is the direct mapping of backup files to Amazon S3, enabling the use of S3 Lifecycle policies.

These policies automatically transition older, less frequently accessed backups to lower-cost storage classes like S3 Glacier (or S3 Glacier Deep Archive) based on defined retention rules, significantly reducing storage costs over time. S3's inherent redundancy across multiple Availability Zones ensures high availability for restoration.

Options A and B are less optimal. Option A using Volume Gateway would involve storing the entire volume data in AWS, which can be more expensive than storing just the backup files. Option B uses Tape Gateway and S3 Glacier Deep Archive directly, which might be too restrictive and not as agile for restoring specific backups when needed. The added step of managing tapes adds complexity. Option C, using third-party imaging and snapshots sent directly to Glacier, is also less cost-effective than using S3 Lifecycle policies to tier storage to Glacier. Snapshot management, especially at 1-hour intervals, leads to significant storage costs. Furthermore, it doesn't utilize the seamless integration and storage tiering capabilities of S3 File Gateway.

Therefore, option D best balances cost-effectiveness, ease of integration, high availability through S3, and long-term retention through S3 Lifecycle policies and archiving to Glacier.

Authoritative Links:

AWS Storage Gateway:https://aws.amazon.com/storagegateway/
Amazon S3 Lifecycle:https://docs.aws.amazon.com/AmazonS3/latest/userguide/lifecycle-configuration-concept.html
Amazon S3 Storage Classes:https://aws.amazon.com/s3/storage-classes/

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: B

Explanation:

The correct answer is **B. Enable versioning and multi-factor authentication (MFA) on the S3 bucket.**

Here's why:

Versioning: S3 Versioning keeps multiple versions of an object in the same bucket. When versioning is enabled, deleting an object doesn't permanently delete it; instead, it creates a delete marker. The previous version remains accessible. This crucial feature protects against accidental deletions, fulfilling the requirement to prevent loss of backup files.
MFA Delete: MFA Delete adds an extra layer of security. To permanently delete a versioned object or change the versioning state of the bucket, the user must provide a valid MFA code from an authorized device. This prevents unauthorized or accidental deletion, specifically addressing the concern raised by the SAP basis team.

Let's analyze why the other options are less suitable:

A. Create a new resource-based policy that prevents deletion of the S3 bucket: While a resource-based policy can restrict access and deletion permissions at the bucket level, it doesn't prevent accidental deletion by authorized users. A user with s3:DeleteBucket permission could still delete the bucket itself, or with sufficient permissions delete the objects within the bucket. It doesn't offer the granular control needed for protecting individual backup versions.

C. Create signed cookies for the backup files in the S3 bucket. Provide the signed cookies to authorized users only: Signed cookies are used to control access to content, especially through CloudFront. They don't prevent deletion. They primarily focus on authorization for viewing or downloading data, not preventing its modification or deletion.

D. Apply an S3 Lifecycle policy to move the backup files immediately to S3 Glacier: While S3 Glacier is cost-effective for long-term archiving, it introduces retrieval costs and delays that are unsuitable for the requirement to restore the last 3 months of backups frequently for testing and performance analysis. Glacier is not designed for frequent restores. Moreover, it doesn't address the accidental deletion problem; if someone with sufficient permissions accidentally triggers a deletion, the objects can still be removed (though it would be a deletion from Glacier, not standard S3).

Therefore, enabling versioning and MFA Delete is the most effective solution to prevent accidental deletion of backup files and ensure the backups are readily available for testing and analysis.

Authoritative Links:

S3 Versioning:https://docs.aws.amazon.com/AmazonS3/latest/userguide/versioning-workflows.html **S3 MFA Delete:**
https://docs.aws.amazon.com/AmazonS3/latest/userguide/Versioning.html#MultiFactorAuthenticationDelete **S3 Bucket Policies:**https://docs.aws.amazon.com/AmazonS3/latest/userguide/bucket-policies.html **S3 Signed Cookies:**https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-signed-cookies.html
S3 Lifecycle Policies:https://docs.aws.amazon.com/AmazonS3/latest/userguide/lifecycle-configuration-examples.html

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? (Choose 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 DR 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 DR 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 DR 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: BE

Explanation:

The question focuses on setting up a highly available and disaster recovery solution for SAP HANA on AWS using system replication. The key constraints are the latency requirements: <1ms within the Availability Zones (AZs) of eu-central-1 and >1ms between eu-central-1 and eu-west-1.

Option B is correct because it places the primary and secondary HANA systems (Tier 1 & Tier 2) within the same region (eu-central-1) but in different AZs. Given the low intra-region latency, synchronous replication (SYNC mode) can be used between these two tiers, ensuring high availability with minimal data loss in case of an AZ failure. The disaster recovery system (Tier 3) in a different region (eu-west-1) necessitates
asynchronous replication (ASYNC mode) due to higher inter-region latency. Synchronous replication across regions would severely impact performance.

Option E is correct because it places the primary and secondary HANA systems (Tier 1 & Tier 2) within the same region (eu-central-1) but in different AZs. The recommendation of using SYNCMEM replication between tier 1 and tier 2 provides the highest level of data consistency by synchronously replicating memory contents.

This is suitable for the low latency environment described. Further, DR system (Tier 3) in a different region (eu-west-1) is achieved using ASYNC mode for reasons mentioned above.

Option A is incorrect because it uses ASYNC between Tier 1 and Tier 2. The latency requirements for the availability zones dictate using a more synchronous method for high availability. Options C and D are incorrect because they either use backups for DR (which has a higher RTO/RPO than system replication), or they improperly place the secondary systems.

Relevant links:

SAP HANA System Replication
SAP on AWS - Disaster Recovery

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 aws-backint-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: B

Explanation:

The problem is long-running database backups using AWS Backint agent. Option B addresses this directly by

suggesting an increase to the UploadChannelSize parameter within the AWS Backint agent configuration. UploadChannelSize controls the number of parallel upload channels to Amazon S3. Increasing this value allows AWS Backint agent to upload data concurrently, potentially significantly reducing the backup time, especially for a large database like 10 TB. This improves throughput by maximizing parallel data transfer to S3.

Option A proposes sending backups over the internet via a NAT gateway. This is generally not recommended for large data transfers due to potential bandwidth limitations, security concerns (exposure to the public internet), and increased cost compared to private network options. The internet introduces network variability and potential bottlenecks. It's also less secure than keeping traffic within the AWS environment.

Option C focuses on MaximumConcurrentFilesForRestore, which only affects restore operations, not backups. Therefore, changing this parameter won't address the long backup times.

Option D suggests compression. While compression can reduce the size of the backup and potentially the time to upload, it's not the most direct solution to the bottleneck. The core issue likely lies in the limited parallelism of the upload process rather than the size of the data itself. Increasing UploadChannelSize will improve the utilization of the network and storage resources.

Therefore, the most effective and direct solution is to increase the UploadChannelSize parameter to improve parallel upload capabilities and reduce backup duration.

Relevant links for further research:

AWS Backint Agent for SAP HANA - Configuration Parameters
AWS Backint Agent for SAP HANA

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 is running on SAP HANA.
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

Explanation:

The correct answer is D. Here's why:

The question focuses on minimizing downtime during the migration while changing the operating system (OS). Different migration strategies have varying downtime implications.

SAP HANA System Replication (HSR) for SAP S/4HANA and SAP BW/4HANA: HSR provides near-zero downtime migration capabilities. Data is continuously replicated from the source system to a target system in AWS. Once replication is complete, a controlled failover can be performed to activate the target system, resulting in minimal downtime. This is ideal for the large databases associated with S/4HANA and BW/4HANA.

SAP HANA System Replication

SAP Software Provisioning Manager (SWPM) for SAP Solution Manager: Solution Manager uses a HANA database, but it's generally smaller in size than S/4HANA or BW/4HANA. An export/import approach using SWPM is suitable for smaller systems where downtime requirements are less stringent. While not as fast as HSR, it's acceptable for this particular component. Changing the OS necessitates this type of procedure.

SAP Software Provisioning Manager

Reinstallation of SAP Web Dispatcher: Web Dispatcher is a relatively lightweight component. It's easier and faster to simply reinstall it on the target system with the necessary configuration. There's no data to migrate, and configuration can be automated.

Let's analyze why the other options are less suitable:

Option A: Export/import for all systems is time-consuming and would lead to significant downtime, which contradicts the question's requirements.

Option B: Backup and restore are also slower than HSR and would result in longer downtime, especially for S/4HANA and BW/4HANA. While suitable for smaller databases, they aren't ideal for the stringent downtime requirement.

Option C: Using backup and restore for SAP S/4HANA and SAP BW/4HANA leads to higher downtime. HSR is a better fit.

In summary, the combination of HSR for the large systems (S/4HANA, BW/4HANA), SWPM export/import for Solution Manager (OS change requirement), and reinstalling Web Dispatcher provides the best balance of minimal downtime and OS migration for this SAP landscape. HSR minimizes downtime for the critical large systems, SWPM handles the OS change for the smaller system, and reinstalling Web Dispatcher is the simplest approach.

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 replicating 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: C

Explanation:

The correct answer is C because it leverages Cross-Platform Transportable Tablespaces (XTTS) for the database migration, a recognized method for heterogeneous database migrations, especially for large databases like the 15 TB Oracle database mentioned in the question.

Here's a breakdown of why C is superior to the other options:

Option A (R3load): R3load is suitable for smaller databases and systems. Given the 15 TB size, a full export and import using R3load during the migration window would be excessively time-consuming, failing to meet the requirement for a fast migration.

Option B (Oracle Data Guard): While Oracle Data Guard is excellent for disaster recovery and high availability, setting it up for a cross-platform migration, especially from IBM Power to AWS, is more complex and might not be the fastest route compared to XTTS, particularly when considering potential compatibility issues. **Option C (XTTS):** XTTS is designed for efficient cross-platform migrations. The pre-migration work of setting up a staging database, performing the initial XTTS conversion, and copying the converted backup to AWS significantly reduces the downtime required during the final cutover window. Incremental XTTS conversions and the transfer of these incremental changes keep the target database relatively up-to-date. During the final migration window, only the final incremental backup and replay of logs are required. This approach minimizes downtime and leverages the 100 Gbps Direct Connect link effectively for data transfer.

Option D (AWS SMS): AWS SMS (Server Migration Service) is more suitable for lift-and-shift migrations of entire servers, not for migrating a database to a different platform. SMS creates AMIs of the source server. Migrating an AMI of an IBM Power server would not allow you to run it directly as a new SAP on Oracle system on the x86_64 architecture of EC2. It would be an inappropriate and ineffective choice. SMS isn't designed for heterogeneous migrations and won't perform the necessary data conversions for an Oracle database from IBM Power to AWS.

In summary: Option C provides the best approach by pre-converting the database using XTTS and transferring the bulk of the data before the migration window. Only the final incremental changes need to be applied during the maintenance window, minimizing downtime and meeting the requirement for the fastest migration.

Authoritative Links:

SAP Note 2632579 - Using Transportable Tablespaces for migrating SAP systems to AWS: https://aws.amazon.com/blogs/awsforsap/migrating-sap-systems-to-aws-using-transportable-tablespaces/ **Oracle Cross-Platform Transportable Tablespaces:**https://docs.oracle.com/en/database/oracle/oracle-database/19/sutil/transporting-tablespaces-between-platforms.html

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 xte.32xiarge Amazon EC2 instances.
The SAP solutions architect must ensure that the SAP HANA scale-out nodes can achieve the low-latency 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? (Choose 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: AD

Explanation:

The correct answer is **AD**. Here's why:

A. Create a cluster placement group. Launch the instances into the cluster placement group.

Cluster placement groups are designed for applications requiring low latency and high network throughput between instances. By launching the EC2 instances into a cluster placement group, you ensure they are placed close together within an AWS Availability Zone. This proximity minimizes latency and maximizes network performance, which is crucial for SAP HANA scale-out architectures where nodes need to communicate frequently and rapidly. Cluster placement groups are ideal for tightly coupled workloads like SAP HANA that depend on high bandwidth and low latency inter-node communication.

D. Based on the operating system version, verify that enhanced networking is enabled on all the nodes.

Enhanced networking utilizes Single Root I/O Virtualization (SR-IOV) to provide higher performance (higher bandwidth, higher packets per second (PPS), and lower latency) on supported instance types. For SAP HANA scale-out on AWS, enabling enhanced networking on all nodes is vital to achieve the required network performance for node-to-node communication. Different operating systems and instance types might require different configurations to enable enhanced networking. Verifying that it's properly configured ensures that the instances can leverage the available network capabilities fully. The x1e.32xlarge instance type supports enhanced networking.

Why the other options are incorrect:

B. Create a spread placement group. Launch the instances into the spread placement group. Spread placement groups aim to distribute instances across distinct underlying hardware, reducing the risk of correlated failures. While good for availability, this distribution can negatively impact network latency and is therefore unsuitable for SAP HANA scale-out, which prioritizes low-latency communication.

C. Create a partition placement group. Launch the instances into the partition placement group. Partition placement groups offer a balance between spread and cluster placement groups. They spread instances across logical partitions within an Availability Zone but do not guarantee the same level of low latency as cluster placement groups. The SAP HANA requirements for high-throughput and low-latency inter-node communication are better served by the proximity provided by cluster placement groups.

E. Switch to a different instance family that provides network throughput that is greater than 25 Gbps. The x1e.32xlarge instance family provides up to 25 Gbps of network performance, which is a suitable baseline. The core issue isn't necessarily the theoretical maximum, but achieving consistently low latency between the nodes. The combination of cluster placement group and enabled enhanced networking is the best approach to meeting the needs. Changing the instance might yield further improvements, but it is not the initial step to address latency.

Supporting Links:

Placement Groups
Enhanced Networking on EC2

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

Explanation:

The correct answer is **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.**

Here's why this is the best solution:

Direct Connect for High Bandwidth, Low Latency: AWS Direct Connect establishes a dedicated network connection from your on-premises environment to AWS. This is crucial for the company's requirement of high bandwidth and low latency, vital for migrating large databases (10 TB+) with minimal downtime. Direct Connect bypasses the public internet, providing a more consistent and predictable network experience.https://aws.amazon.com/directconnect/

Site-to-Site VPN for Resiliency: While Direct Connect is the primary connection, AWS Site-to-Site VPN provides a backup connection over the internet. In case of Direct Connect failure, the VPN connection ensures continued connectivity, albeit at a lower bandwidth and higher latency. This fulfills the requirement for connectivity resiliency.https://aws.amazon.com/vpn/

Why other options are not optimal:

B. Direct Connect Gateway with LAG: While LAG improves bandwidth, it doesn't inherently provide redundancy in the event of a complete Direct Connect outage at the physical connection level. It's more for increased throughput.

C. Amazon EFS: EFS is a shared file system, not designed for large-scale database migrations. Copying data to EFS via a cron job introduces significant overhead and would be far slower than Direct Connect. Also, transferring large databases to EFS is not a standard migration strategy and can be costly.https://aws.amazon.com/efs/

D. Redundant Site-to-Site VPN: VPN connections are internet-based and don't provide the consistent high

bandwidth and low latency that Direct Connect offers. Although redundant VPNs improve availability compared to a single VPN, they are unsuitable for minimizing downtime during large database migrations.

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? (Choose 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 route on the production VPC to route traffic that is directed to the overlay IP address to the ASCS instance.

Answer: BD

Explanation:

The correct answer is BD. Let's break down why:

Understanding the Requirements: The company needs highly available access to SAP systems in a production VPC from their on-premises data center via a Site-to-Site VPN. Crucially, we're focusing on how to manage the ASCS (ABAP SAP Central Services) instance for high availability.

Option B: Overlay IP Outside VPC CIDR with Transit Gateway: The ASCS instance typically uses an overlay IP address for high availability. Configuring this IP address outside the VPC CIDR range is preferable. This ensures it doesn't conflict with any existing or future IP address assignments within the VPC. The AWS Transit Gateway acts as a central hub for routing traffic between the on-premises data center (via the VPN) and the production VPC. The Transit Gateway simplifies network management significantly. A static route in the production VPC ensures traffic destined for the ASCS overlay IP gets routed correctly to the ASCS instance.

Option D: Overlay IP Outside VPC CIDR with NLB: As in option B, using an overlay IP outside the VPC CIDR range is a good practice to avoid conflicts. The Network Load Balancer (NLB) is the correct load balancer choice for SAP workloads, as it preserves the source IP address of the client, which is crucial for SAP authentication and authorization. The target group associated with the NLB directs traffic to the ASCS instance's overlay IP address. A static route in the production VPC ensures that any traffic to the ASCS overlay IP gets directed to the NLB.

Why other options are incorrect:

Option A and C: Using an overlay IP address within the VPC CIDR range is a bad practice. It can lead to IP address conflicts and routing issues.

Option E: The Application Load Balancer (ALB) is not suitable for SAP workloads. The ALB modifies the source IP address, which breaks critical SAP functionality. NLB is the recommended choice for SAP.

Therefore, options B and D provide configurations that correctly handle high availability for the ASCS instance while ensuring seamless connectivity between the on-premises data center and the production VPC.

Supporting links:

AWS Transit Gateway: https://aws.amazon.com/transit-gateway/
Network Load Balancer: https://aws.amazon.com/elasticloadbalancing/network-load-balancer/ SAP on AWS Implementation Guide: https://aws.amazon.com/sap/

Question: 14

A company is running an SAP HANA database on AWS. The company is running AWS Backint Agent for SAP HANA (AWS Backint agent) on an Amazon EC2 instance. AWS Backint agent is configured to back up to an Amazon S3 bucket. The backups are failing with an AccessDenied error in the AWS Backint agent log file.
What should an SAP basis administrator do to resolve this error?

A.Assign execute permissions at the operating system level for the AWS Backint agent binary and for AWS Backint agent.

B.Assign an IAM role to an EC2 instance. Attach a policy to the IAM role to grant access to the target S3 bucket.

C.Assign the correct Region ID for the S3BucketAwsRegion parameter in AWS Backint agent for the SAP HANA configuration file.

D.Assign the value for the EnableTagging parameter in AWS Backint agent for the SAP HANA configuration file.

Answer: B

Explanation:

The correct answer is B: Assign an IAM role to the EC2 instance and attach a policy granting access to the target S3 bucket. The AccessDenied error indicates that the AWS Backint agent, running on the EC2 instance, lacks the necessary permissions to write backups to the designated S3 bucket. AWS best practices emphasize granting permissions to AWS resources using IAM roles.

IAM roles provide a secure way to grant permissions to applications running on EC2 instances, without needing to manage AWS credentials directly within the instance. This eliminates the need to store sensitive access keys and secret keys on the EC2 instance itself, improving security posture.

By assigning an IAM role to the EC2 instance, and then attaching an IAM policy to that role that allows s3:PutObject (and potentially other necessary S3 actions) to the specific S3 bucket, the AWS Backint agent will inherit the necessary permissions to perform backups. Option A is incorrect because execute permissions are already required for the binary to run. Option C addresses an incorrect region config, which may cause a different error. Option D is related to tagging, not access control.

Relevant AWS documentation:

IAM Roles for EC2: https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_ec2.html IAM Policies: https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies.html
S3 Permissions: https://docs.aws.amazon.com/AmazonS3/latest/userguide/using-with-s3-actions.html

Question: 15

A company is starting a new project to implement an SAP landscape with multiple accounts that belong to multiple teams in the us-east-2 Region. These teams include procurement, finance, sales, and human resources. An SAP solutions architect has started designing this new landscape and the AWS account structures.
The company wants to use automation as much as possible. The company also wants to secure the environment, implement federated access to accounts, centralize logging, and establish cross-account security audits. In addition, the company's management team needs to receive a top-level summary of policies that are applied to the AWS accounts.
What should the SAP solutions architect do to meet these requirements?

A.Use AWS CloudFormation StackSets to apply SCPs to multiple accounts in multiple Regions. Use an Amazon CloudWatch dashboard to check the applied policies in the accounts.

B.Use an AWS Elastic Beanstalk blue/green deployment to create IAM policies and apply them to multiple accounts together. Use an Amazon CloudWatch dashboard to check the applied policies in the accounts.

C.Implement guardrails by using AWS CodeDeploy and AWS CodePipeline to deploy SCPs into each account. Use the CodePipeline deployment dashboard to check the applied policies in the accounts.

D.Apply SCPs through AWS Control Tower. Use the AWS Control Tower integrated dashboard to check the applied policies in the accounts.

Answer: D

Explanation:

The correct answer is **D: Apply SCPs through AWS Control Tower. Use the AWS Control Tower integrated dashboard to check the applied policies in the accounts.**

Here's why:

AWS Control Tower is designed precisely for setting up and governing multi-account AWS environments, especially aligning with organizational best practices. It automates the creation of landing zones (multi-account structures) and helps enforce consistent policies across those accounts using Service Control Policies (SCPs). The question highlights key requirements like automation, security, federated access, centralized logging, cross-account security audits, and top-level policy summaries. Control Tower directly addresses these.

Automation & Multi-Account Management: Control Tower automates the setup of a multi-account environment, provisioning new accounts and applying baseline configurations.

Security: SCPs, applied via Control Tower, enable centralized control over the permissions available in each account. They act as guardrails, preventing actions even by administrators within individual accounts if those actions violate the centrally defined policies.

Federated Access: Control Tower integrates with AWS IAM Identity Center (successor to AWS SSO) to provide federated access to accounts using existing identities.

Centralized Logging & Auditing: Control Tower centralizes logs and audit information into a central account for compliance and security analysis.

Policy Summary: The Control Tower dashboard provides a high-level overview of the policies applied across the organization.

The other options are not ideal:

A (CloudFormation StackSets): While StackSets can deploy SCPs, they lack the comprehensive governance and management features of Control Tower. There is no consolidated dashboard for checking applied policies.

B (Elastic Beanstalk): Elastic Beanstalk is for application deployment, not for managing IAM policies across multiple accounts. It's fundamentally unrelated to the requirements.

C (CodeDeploy & CodePipeline): CodeDeploy and CodePipeline are CI/CD tools for deploying applications and infrastructure changes. While they can automate SCP deployment, they do not offer the centralized management, auditing, and dashboarding capabilities of Control Tower. Also, using CodePipeline for SCPs

lacks the built-in guardrail features Control Tower provides.

In summary, AWS Control Tower provides the most comprehensive and purpose-built solution to meet all the requirements outlined in the question, making it the best choice.

Relevant links for further research:

AWS Control Tower: https://aws.amazon.com/controltower/
Service Control Policies (SCPs):
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html AWS
IAM Identity Center (successor to AWS SSO): https://aws.amazon.com/iam/identity/

Question: 16

A company is running its SAP workloads on premises and needs to migrate the workloads to AWS. All the workloads are running on SUSE Linux Enterprise Server and Oracle Database. The company's landscape consists of SAP ERP Central Component (SAP ECC), SAP Business Warehouse (SAP BW), and SAP NetWeaver systems. The company has a dedicated AWS Direct Connect connection between its on-premises environment and AWS. The company needs to migrate the systems to AWS with the least possible downtime.
Which migration solution will meet these requirements?

A.Use SAP Software Provisioning Manager to perform an export of the systems. Copy the export to Amazon S3. Use SAP Software Provisioning Manager to perform an import of the systems to SUSE Linux Enterprise Server and Oracle Database on AWS.

B.Use SAP Software Provisioning Manager to perform parallel export/import of the systems to migrate the systems to SUSE Linux Enterprise Server and Oracle Database on AWS.

C.Use SAP Software Provisioning Manager to perform parallel export/import of the systems to migrate the systems to Oracle Enterprise Linux and Oracle Database on AWS.

D.Use SAP Software Provisioning Manager to perform an export of the systems. Copy the export to Amazon S3. Use SAP Software Provisioning Manager to perform an import of the systems to Oracle Enterprise Linux and Oracle Database on AWS.

Answer: C

Explanation:

The correct answer is C because it leverages the capabilities of SAP Software Provisioning Manager (SWPM) for a heterogeneous system copy, aiming for minimal downtime during the migration to AWS.

Here's a breakdown of why C is preferred and why the others are less suitable:

Heterogeneous System Copy with SWPM: The company is moving from an on-premises environment to AWS, which involves a platform change. The supported OS and database are changing to Oracle Enterprise Linux (OEL) and Oracle Database, respectively. SWPM is designed to handle such heterogeneous migrations, involving changes in OS and database. It facilitates the migration by exporting the existing SAP data and importing it into the new environment.
https://help.sap.com/docs/SOFTWARE_PROVISIONING_MGR/251326b076854488b710b7168880b368/987a01753

Parallel Export/Import: SWPM allows parallel execution of export and import processes, significantly reducing downtime during the migration. This approach aligns with the requirement of minimizing downtime.

Oracle Enterprise Linux and Oracle Database on AWS: SAP on AWS best practices often recommend using Oracle Enterprise Linux and Oracle Database due to performance optimization and support.

Why other options are incorrect:

A & B: Maintaining SUSE Linux Enterprise Server and Oracle Database in AWS might be a straightforward lift

and shift migration. However, the option is not using the recommended OS and DB for SAP on AWS. While technically feasible, staying with SUSE might not be the optimal choice for performance or cost depending on the specific workload and AWS offerings.

D: While using Oracle Enterprise Linux and Oracle Database is preferred, the non-parallel approach will lead to increased downtime. Copying the export to S3 adds an extra step that will increase the total migration time.

Question: 17

A company is designing a disaster recovery (DR) strategy for an SAP HANA database that runs on an Amazon EC2 instance in a single Availability Zone. The company can tolerate a long RTO and an RPO greater than zero if it means that the company can save money on its DR process.
The company has configured an Amazon CloudWatch alarm to automatically recover the EC2 instance if the instance experiences an unexpected issue. The company has set up AWS Backint Agent for SAP HANA to save the backups into Amazon S3.
What is the MOST cost-effective DR option for the company's SAP HANA database?

A.Set up AWS CloudFormation to automatically launch a new EC2 instance for the SAP HANA database in a second Availability Zone from backups that are stored in Amazon S3. When the SAP HANA database is operational, perform a database restore by using the standard SAP HANA restore process.

B.Launch a secondary EC2 instance for the SAP HANA database on a less powerful EC2 instance type in a second Availability Zone. Configure SAP HANA system replication with the preload option turned off.

C.Launch a secondary EC2 instance for the SAP HANA database on an equivalent EC2 instance type in a second Availability Zone. Configure SAP HANA system replication with the preload option turned on.

D.Set up AWS CloudFormation to automatically launch a new EC2 instance for the SAP HANA database in a second Availability Zone from backups that are stored in Amazon Elastic Block Store (Amazon EBS). When the SAP HANA database is operational, perform a database restore by using the standard SAP HANA restore process.

Answer: A

Explanation:

The correct answer is A because it offers the most cost-effective disaster recovery solution that aligns with the company's tolerance for a longer RTO and an RPO greater than zero.

Here's why:

Cost-Effectiveness: Launching a new EC2 instance using CloudFormation only when needed avoids the cost of running a standby instance continuously. Storing backups in S3 is also cheaper than EBS for long-term storage.

Meets RTO/RPO Requirements: The company is okay with a longer recovery time objective (RTO). Restoring from S3 backups will take time, but it's acceptable. The RPO is also satisfied as backups are taken, allowing for some data loss.

CloudFormation Automation: CloudFormation simplifies and automates the recovery process, ensuring consistency and reducing manual intervention, which is a significant advantage for disaster recovery scenarios.

Less Powerful Instance (Option B): Although Option B (less powerful instance with SAP HANA System Replication and preload off) would reduce costs, it's still more expensive than launching an instance only when needed and restoring from backups. The continuous operation of even a smaller instance incurs costs. Moreover, restoring from S3 may have additional latency involved with the S3 endpoint.

Options C and D are incorrect:

Option C (SAP HANA System Replication with preload on): This is the most expensive option because it requires a fully sized, identical EC2 instance running continuously. It is also not cost-effective.

Option D (Restoring from EBS): This is potentially more expensive than S3 for long-term storage and doesn't provide a substantial advantage in terms of recovery time, given the acceptable RTO.

Here are some useful links for further research:

AWS Disaster Recovery: https://aws.amazon.com/disaster-recovery/
AWS CloudFormation: https://aws.amazon.com/cloudformation/
Amazon S3 Storage Classes: https://aws.amazon.com/s3/storage-classes/ AWS Backint
Agent for SAP HANA: https://aws.amazon.com/sap/solutions/hana/

Question: 18

A company is using a multi-account strategy for SAP HANA and SAP BW/4HANA instances across development, QA, and production systems in the same AWS Region. Each system is hosted in its own VPC. The company needs to establish cross-VPC communication between the SAP systems.
The company might add more SAP systems in the future. The company must create connectivity across the SAP systems and hundreds of AWS accounts. The solution must maximize scalability and reliability.
Which solution will meet these requirements?

A.Create an AWS Transit Gateway in a central networking account. Attach the transit gateway to the AWS accounts. Set up routing and a network ACL to establish communication.

B.Set up VPC peering between the accounts. Configure routing in each VPC to use the VPC peering links.

C.Create a transit VPC that uses the hub-and-spoke model. Set up routing to use the transit VPC for communication between the SAP systems.

D.Create a VPC link for each SAP system. Use the VPC links to connect the SAP systems.

Answer: A

Explanation:

The best solution is A: Create an AWS Transit Gateway in a central networking account, attach it to the AWS accounts, and set up routing and network ACLs. Here's why:

Scalability and Centralized Management: AWS Transit Gateway is designed for connecting thousands of VPCs and on-premises networks in a hub-and-spoke topology. Centralizing this in a networking account simplifies management and scaling for future SAP systems and AWS accounts.

Simplified Routing: Transit Gateway simplifies routing configurations. You configure routes within the Transit Gateway route tables to determine how traffic flows between attached VPCs, removing the need for individual VPC peering route configurations.

Reliability: AWS Transit Gateway is a highly available and resilient service, offering high availability and automatic scaling.

Avoids Complexity of Peering: VPC peering (Option B) becomes unmanageable with hundreds of accounts and VPCs. Each new VPC requires additional peering connections, leading to a complex mesh network. **Transit VPC Limitations:** A transit VPC (Option C) is a valid older method, but Transit Gateway provides better scalability, higher bandwidth, and simplified management. Transit VPCs also typically require more configuration and maintenance of network appliances within the transit VPC.

VPC Links with API Gateway: VPC Links (Option D) are for integrating API Gateway with resources in a VPC. They aren't designed for general-purpose cross-VPC networking between SAP systems.

Option A offers a scalable, reliable, and manageable solution for connecting multiple SAP systems across many AWS accounts, which is the most suitable approach based on the prompt's requirements.

Authoritative Links:

AWS Transit Gateway: https://aws.amazon.com/transit-gateway/

AWS Transit Gateway documentation: https://docs.aws.amazon.com/transit-gateway/index.html VPC Peering: https://docs.aws.amazon.com/vpc/latest/peering/what-is-vpc-peering.html

Question: 19

A company is planning to deploy a new SAP NetWeaver ABAP system on AWS with an Oracle database that runs on an Amazon EC2 instance. The EC2 instance uses a Linux-based operating system. The company needs a database storage solution that provides flexibility to adjust the IOPS regardless of the allocated storage size. Which solution will meet these requirements MOST cost-effectively?

A.General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volumes

B.Amazon Elastic File System (Amazon EFS) Standard-Iinfrequent Access (Standard-IA) storage class

C.Amazon FSx for Windows File Server

D.Provisioned IOPS SSD (io2) Amazon Elastic Block Store (Amazon EBS) volumes

Answer: A

Explanation:

The correct answer is **A. General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volumes**. Here's why:

The question emphasizes the need for a cost-effective storage solution that allows independent adjustment of IOPS irrespective of the storage size. This is crucial for SAP systems to handle varying workloads efficiently.

gp3 volumes offer the ability to provision IOPS independently of storage size. This is a key feature that allows the company to scale IOPS as needed without having to increase storage capacity unnecessarily, making it a cost-effective option. gp3 provides a baseline performance and allows you to provision additional IOPS to meet higher performance requirements, and do so cost effectively

io2 volumes, while offering the capability to provision IOPS, are generally more expensive than gp3 volumes. They are intended for the most demanding, I/O intensive workloads, and for the described use case, gp3 is more than adequate.

Amazon EFS, is a network file system, suitable for shared file storage across multiple instances. It's not ideal for database storage, which requires block storage for optimal performance. Additionally, although EFS Performance modes are available that are more suitable for DB workloads (as opposed to the cheaper throughput optimized, IA versions), it is more expensive and not the correct architecture for DBs.

Amazon FSx for Windows File Server is designed for Windows-based applications, not Linux-based Oracle databases. It's also a shared file system and not suitable for the block storage requirements of a database.

In summary, gp3 EBS volumes provide the best balance of performance, flexibility (with adjustable IOPS), and cost-effectiveness for an SAP NetWeaver ABAP system with an Oracle database running on EC2 with a Linux OS. It offers the ability to tailor IOPS independently of storage size, which addresses the explicit requirement of the question, whilst also keeping costs down compared to IO2.

Supporting documentation:

Amazon EBS volume types: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volume-types.html gp3 volumes: https://aws.amazon.com/ebs/general-purpose/

Question: 20

A company is using SAP NetWeaver with Java on AWS. The company has updated its generation of Amazon EC2 instances to the most recent generation of EC2 instances. When the company tries to start SAP, the startup fails. The log indicates that the SAP license expired or is not valid.
What is the reason for this issue?

A.The instance ID changed as part of the EC2 generation change.

B.The instance's hypervisor changed from Xen to Nitro.

C.The SAP Java Virtual Machine (SAP JVM) is not compatible with the new instance type.

D.An EC2 generation change is not supported for SAP Java-based systems.

Answer: B

Explanation:

The correct answer is B: The instance's hypervisor changed from Xen to Nitro.

Here's a detailed justification:

SAP licensing often relies on hardware attributes of the underlying infrastructure. One key attribute that's often used in SAP licensing is tied to the server's hardware key, which is derived from characteristics reported by the hypervisor. The transition from older EC2 instance generations to newer generations often involves a move from the Xen hypervisor to the Nitro hypervisor system.

This hypervisor change alters the underlying hardware characteristics exposed to the operating system and, consequently, to SAP. SAP systems, particularly older implementations, might not recognize the new hardware key derived from the Nitro hypervisor as valid against the existing license tied to the Xen-based hardware key. The SAP system then interprets this as an expired or invalid license.

Option A is incorrect because while the instance ID might change during certain migrations (e.g., creating a new instance from an AMI), a simple EC2 instance generation change does not automatically alter the instance ID. The license is typically not based on the EC2 instance ID.

Option C is incorrect because SAP JVM compatibility is generally tied to the operating system and SAP version, not directly to the EC2 instance type. While some compatibility issues could theoretically arise in extreme cases with very old JVM versions, it is highly unlikely to be the root cause of the license invalidation described. SAP generally ensures compatibility between the JVM and standard hardware.

Option D is incorrect because EC2 instance generation changes are, in general, supported for SAP systems. However, the licensing impact must be considered and managed. The licensing changes from Xen to Nitro are possible when upgrading from a previous generation instance to a Nitro-based instance.

Therefore, the most likely reason for the SAP startup failure with a license error after an EC2 instance generation change is the change in the underlying hypervisor, specifically from Xen to Nitro, affecting the hardware key used for SAP licensing.

Further Research:

AWS Nitro System:https://aws.amazon.com/ec2/nitro/
SAP Licensing: SAP licensing information is available at https://www.sap.com/ (look for licensing FAQs and documentation)

Question: 21

A company's basis administrator is planning to deploy SAP on AWS in 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? (Choose two.)

A.Amazon Elastic Block Store (Amazon EBS) Throughput Optimized HDD (st1)

B.Amazon Elastic Block Store (Amazon EBS) Provisioned OPS SSD (io1, io2)

C.Amazon S3

D.Amazon Elastic File System (Amazon EFS)

E.Amazon Elastic Block Store (Amazon EBS) General Purpose SSD (gp2, gp3)

Answer: BE

Explanation:

The correct answer is BE. Here's a detailed justification:

Amazon EBS Provisioned IOPS SSD (io1, io2): SAP HANA data and log volumes require high performance and low latency. Provisioned IOPS SSD (io1, io2) EBS volumes are designed for I/O-intensive workloads, allowing the basis administrator to specify the required IOPS (Input/Output Operations Per Second) to meet the SAP HANA performance needs. This is crucial for the responsiveness and efficiency of the SAP system.
https://aws.amazon.com/ebs/

Amazon EBS General Purpose SSD (gp2, gp3): gp2 and gp3 EBS volumes are general-purpose SSD volumes that provide a balance of price and performance suitable for a wide variety of workloads. gp3 offers greater flexibility than gp2 since it allows the ability to independently scale IOPS and throughput without having to provision additional block storage capacity. They can be used for SAP HANA data and log volumes but require careful sizing and monitoring to ensure performance requirements are met.
https://aws.amazon.com/ebs/general-purpose/

Let's analyze why the other options are not suitable:

Amazon Elastic Block Store (Amazon EBS) Throughput Optimized HDD (st1): st1 volumes are designed for frequently accessed, throughput-intensive workloads with large datasets and large I/O sizes, such as big data, data warehouses, and log processing. They are not suitable for the low-latency, high-IOPS requirements of SAP HANA data and log volumes.

Amazon S3: Amazon S3 is object storage and is not designed for use as a file system for transactional databases like SAP HANA. SAP HANA requires block storage volumes.

Amazon Elastic File System (Amazon EFS): Amazon EFS provides a scalable, elastic, shared file system. While it can be used with Linux-based applications, it generally does not provide the low latency and high IOPS required for SAP HANA data and log volumes, especially for production environments.

Question: 22

A company has deployed a highly available SAP NetWeaver system on SAP HANA into a VPC. The system is distributed across multiple Availability Zones within a single AWS Region. SAP NetWeaver is running on SUSE Linux Enterprise Server for SAP. SUSE Linux Enterprise High Availability Extension is configured to protect SAP ASCS and ERS instances and uses the overlay IP address concept. The SAP shared files /sapmnt and
/usr/sap/trans are hosted on an Amazon Elastic File System (Amazon EFS) file system.
The company needs a solution that uses already-existing private connectivity to the VPC. The SAP NetWeaver system must be accessible through the SAP GUI client tool.
Which solutions will meet these requirements? (Choose two.)

A.Deploy an Application Load Balancer. Configure the overlay IP address as a target.

B.Deploy a Network Load Balancer. Configure the overlay IP address as a target.

C.Use an Amazon Route 53 private zone. Create an A record that has the overlay IP address as a target.

D.Use AWS Transit Gateway. Configure the overlay IP address as a static route in the transit gateway route table. Specify the VPC as a target.

E.Use a NAT gateway. Configure the overlay IP address as a target.

Answer: BD

Explanation:

Let's break down why options B and D are correct, and why the others are not, in the context of SAP NetWeaver HA on AWS with overlay IP addresses.

B. Deploy a Network Load Balancer. Configure the overlay IP address as a target.

Justification: The key here is the "overlay IP address" used by SUSE's High Availability Extension for SAP ASCS/ERS. This IP address floats between the ASCS/ERS instances, providing a single, consistent entry point for SAP applications, even if the underlying instance fails. A Network Load Balancer (NLB) is designed to handle static IP addresses and TCP/UDP traffic, making it the ideal choice for directing SAP GUI client connections to the active ASCS/ERS instance via the overlay IP. NLBs preserve the source IP address of the client, which can be important for SAP security and auditing. NLBs operate at Layer 4, forwarding traffic directly to the target IP address and port. It offers high performance and low latency which is crucial for SAP applications.

Relevance to Question: The question mentions the SAP NetWeaver system must be accessible via SAP GUI. The NLB ensures this requirement by enabling the SAP GUI client to connect to the floating IP, directing the traffic to the active ASCS/ERS instance.

Why it works: NLB directly associates with IP addresses as targets. HA cluster manages which server "owns" this address and NLB just forwards accordingly.

D. Use AWS Transit Gateway. Configure the overlay IP address as a static route in the transit gateway route table. Specify the VPC as a target.

Justification: The company states they need to use already-existing private connectivity to the VPC. AWS Transit Gateway is a hub-and-spoke service that allows you to connect multiple VPCs and on-premises networks through a central hub. In this scenario, the on-premises SAP GUI clients need to reach the ASCS/ERS instances within the VPC. The overlay IP address is the destination. Adding a static route in the Transit Gateway route table, with the overlay IP as the destination and the VPC as the target, ensures that traffic destined for the overlay IP is correctly routed from the on-premises network (or other connected VPCs) to the VPC hosting the SAP system. This meets the requirement of using existing private connectivity. **Relevance to Question:** The question mentions the need for already-existing private connectivity, and Transit Gateway excels at this.

Why it works: Transit Gateway acts as a central routing point and allows to specify static routes.

Why other Options are incorrect:

A. Deploy an Application Load Balancer. Configure the overlay IP address as a target. Application Load Balancers (ALBs) are designed for HTTP/HTTPS traffic and operate at Layer 7. While ALBs can provide advanced routing based on content, they aren't the right choice for SAP GUI clients, which typically use a proprietary SAP protocol over TCP. ALBs are not suited for forwarding traffic directly to IP addresses. ALBs typically require an HTTP/HTTPS application running on the target.

C. Use an Amazon Route 53 private zone. Create an A record that has the overlay IP address as a target.

Route 53 is a DNS service. While you can use it within a VPC, it resolves domain names to IP addresses. It doesn't provide routing or traffic management capabilities in the same way as a load balancer or Transit Gateway. Simply resolving the domain name to the overlay IP doesn't guarantee that traffic will be directed to the active ASCS/ERS instance. DNS is only consulted for address resolution, not for directing traffic based on availability.

E. Use a NAT gateway. Configure the overlay IP address as a target. NAT Gateways are for enabling instances in private subnets to connect to the internet or other AWS services. They are not used for routing traffic to a specific IP address within a VPC, especially not an overlay IP address that is managed by a high

availability cluster. NAT Gateways are unidirectional (outbound only).

Authoritative Links:

Network Load Balancer:https://docs.aws.amazon.com/elasticloadbalancing/latest/network/introduction.html **AWS Transit Gateway:**https://docs.aws.amazon.com/transit-gateway/latest/tgw/what-is-transit-gateway.html
SUSE Linux Enterprise High Availability Extension: (Check the SAP documentation or SUSE's website)

Question: 23

A company is planning to move all its SAP applications to Amazon EC2 instances in a VPC. Recently, the company signed a multiyear contract with a payroll software-as-a-service (SaaS) provider. Integration with the payroll SaaS solution is available only through public web APIs.
Corporate security guidelines state that all outbound traffic must be validated against an allow list. The payroll SaaS provider provides only fully qualified domain name (FQDN) addresses and no IP addresses or IP address ranges. Currently, an on-premises firewall appliance filters FQDNs. The company needs to connect an SAP Process Orchestration (SAP PO) system to the payroll SaaS provider.
What must the company do on AWS to meet these requirements?

A.Add an outbound rule to the security group of the SAP PO system to allow the FQDN of the payroll SaaS provider and deny all other outbound traffic.

B.Add an outbound rule to the network ACL of the subnet that contains the SAP PO system to allow the FQDN of the payroll SaaS provider and deny all other outbound traffic.

C.Add an AWS WAF web ACL to the VPAdd an outbound rule to allow the SAP PO system to connect to the FQDN of the payroll SaaS provider.

D.Add an AWS Network Firewall firewall to the VPC. Add an outbound rule to allow the SAP PO system to connect to the FQDN of the payroll SaaS provider.

Answer: D

Explanation:

The correct answer is D: Add an AWS Network Firewall firewall to the VPC. Add an outbound rule to allow the SAP PO system to connect to the FQDN of the payroll SaaS provider.

Here's why:

Requirement for FQDN Filtering: The company requires filtering outbound traffic based on FQDNs, not just IP addresses. Traditional security groups and network ACLs (NACLs) in AWS operate primarily on IP addresses and ports, not FQDNs. AWS WAF is for web application layer protection (HTTP/HTTPS traffic) and isn't suitable for general outbound traffic filtering based on FQDNs.

AWS Network Firewall Capability: AWS Network Firewall is a managed service that allows you to implement fine-grained network traffic filtering based on domain names (FQDNs) using Suricata compatible rules. This directly addresses the requirement.

Outbound Traffic Control: The company needs to control outbound traffic from the SAP PO system to the payroll SaaS provider. Network Firewall sits in the path of network traffic leaving the VPC and can enforce policies accordingly.

Allow List Validation: The solution needs to validate traffic against an allow list of FQDNs. Network Firewall allows configuring rules that permit traffic to specific FQDNs and deny all other outbound traffic, thereby adhering to corporate security guidelines.

Why other options are incorrect:

A: Security Groups: Security groups are stateful firewalls that control inbound and outbound traffic at the instance level. They operate on IP addresses and ports, not FQDNs.

B: Network ACLs: Network ACLs are stateless firewalls that control inbound and outbound traffic at the subnet level. Like security groups, they operate on IP addresses and ports, not FQDNs.

C: AWS WAF: AWS WAF protects web applications from common web exploits and attacks. It is designed for HTTP(S) traffic and cannot control general outbound traffic based on FQDNs. WAF is used in conjunction with services like Application Load Balancer or API Gateway and is not meant to be a general-purpose firewall for VPC outbound traffic.

Authoritative Links:

AWS Network Firewall:https://aws.amazon.com/network-firewall/
AWS Network Firewall Documentation:https://docs.aws.amazon.com/network-firewall/latest/developerguide/what-is-aws-network-firewall.html

Question: 24

A company is planning to migrate its on-premises SAP application to AWS. The application runs on VMware vSphere. The SAP ERP Central Component (SAP ECC) server runs on an IBM Db2 database that is 2 TB in size. The company wants to migrate the database to SAP HANA.
Which migration strategy will meet these requirements?

A.Use AWS Application Migration Service (CloudEndure Migration).
B.Use SAP Software Update Manager (SUM) Database Migration Option (DMO) with System Move.
C.Use AWS Server Migration Service (AWS SMS).
D.Use AWS Database Migration Service (AWS DMS).

Answer: B

Explanation:

The correct answer is **B. Use SAP Software Update Manager (SUM) Database Migration Option (DMO) with System Move.**

Here's why:

Requirement for SAP HANA Migration: The company specifically wants to move their IBM Db2 database to SAP HANA. This is a database platform change, not just a simple lift-and-shift migration.

SUM DMO with System Move: SAP's SUM DMO with System Move is the recommended and supported tool for performing a database migration and an SAP system upgrade/migration simultaneously. It handles the entire process, including schema conversion, data transfer, and SAP system adaptation to the new database. **AWS Application Migration Service (CloudEndure Migration):** CloudEndure, now AWS Application Migration Service, excels at lift-and-shift migrations of entire servers. While it can migrate the SAP application, it doesn't facilitate the Db2 to HANA database conversion.

AWS Server Migration Service (AWS SMS): AWS SMS is another tool for migrating on-premises virtual machines to AWS. It also doesn't handle the database conversion aspect needed for this scenario.

AWS Database Migration Service (AWS DMS): DMS is a good choice for migrating databases, but it typically requires downtime for the switchover or complex configurations for near-zero downtime. SUM DMO is specifically designed for SAP environments and integrates tightly with SAP's upgrade processes, minimizing downtime and complexity within the SAP landscape. In addition, DMS is more suited for homogeneous or heterogeneous database migrations, but it does not provide the specific functionality to convert a database to SAP HANA and adapt the SAP system accordingly like SUM DMO does.

In conclusion, SUM DMO with System Move directly addresses the need to migrate the Db2 database to SAP HANA and adapt the SAP system accordingly, making it the most suitable choice.

SAP SUM DMO Documentation

AWS Application Migration Service
AWS Database Migration Service

Question: 25

A company hosts multiple SAP applications on Amazon EC2 instances in a VPC. While monitoring the environment, the company notices that multiple port scans are attempting to connect to SAP portals inside the VPC. These port scans are originating from the same IP address block. The company must deny access to the VPC from all the offending IP addresses for the next 24 hours.
Which solution will meet this requirement?

A.Modify network ACLs that are associated with all public subnets in the VPC to deny access from the IP address block.

B.Add a rule in the security group of the EC2 instances to deny access from the IP address block.

C.Create a policy in AWS Identity and Access Management (IAM) to deny access from the IP address block.

D.Configure the firewall in the operating system of the EC2 instances to deny access from the IP address block.

Answer: A

Explanation:

The correct answer is A: Modify network ACLs (NACLs) that are associated with all public subnets in the VPC to deny access from the IP address block. Here's why:

Network ACLs (NACLs) provide stateless packet filtering at the subnet level: NACLs act as a firewall for controlling traffic in and out of one or more subnets. They evaluate traffic entering and exiting a subnet, allowing or denying traffic based on defined rules. This makes them ideal for blocking unwanted traffic sources based on IP addresses. https://docs.aws.amazon.com/vpc/latest/userguide/vpc-network-acls.html

Security Groups operate at the instance level: Security Groups act as a virtual firewall for your EC2 instances to control inbound and outbound traffic. While effective, they only protect specific instances, not the entire subnet. This would require modifying each EC2 instance's security group individually, which is less efficient.

IAM Policies control AWS API access, not network traffic: IAM policies govern access to AWS services and resources via the AWS API. They are not designed to block network traffic based on IP addresses.

Operating system firewalls are instance-specific: Configuring the firewall on each EC2 instance's operating system is also possible, but it's less centralized and more complex to manage across multiple instances than using NACLs.

Public subnets are where internet-facing resources reside: Since the port scans are originating from outside the VPC, the attacking traffic is likely entering the VPC through the public subnets, which contain internet-facing resources such as load balancers or NAT gateways.

NACLs are stateless: NACLs perform stateless packet filtering, meaning that they don't keep track of connections. Both inbound and outbound rules must be explicitly defined.

By modifying NACLs associated with the public subnets, the company can efficiently block the offending IP address block from accessing the SAP portals within the VPC for the required 24-hour period without having to modify each EC2 instance individually. This is the most direct and effective approach to meet the stated requirements. The temporary nature of the block (24 hours) makes NACL management an acceptable overhead.

Question: 26

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:

Add the cloudwatch:GetMetricStatistics action to the policy statement with Sid AWSDataProvider1.

Reference:

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

Question: 27

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:

The correct answer is CE, providing the Amazon S3 URL of the script and uploading the script directly.

Here's why: AWS Launch Wizard for SAP enables automated deployments of SAP workloads on AWS. A critical part of automation is the ability to execute post-deployment scripts for customized configurations. Launch Wizard allows flexibility in specifying how these scripts are delivered.

Option C, providing the Amazon S3 URL, is valid because AWS services seamlessly integrate with S3 for object storage and retrieval. S3 offers a secure and scalable method to store and distribute the script to the EC2 instance provisioned by Launch Wizard. The EC2 instance needs appropriate IAM permissions to access the S3 bucket.
https://docs.aws.amazon.com/AmazonS3/latest/userguide/

Option E, uploading the script, is also a valid method. Launch Wizard provides an interface that allows directly uploading the script during the deployment configuration. This eliminates the need for an external server.

Option A, providing an FTP URL, is incorrect. FTP is not a secure protocol, and AWS typically discourages its use, especially for sensitive operations like post-deployment configuration. AWS emphasizes secure communication channels.

Option B, providing an HTTPS URL on a web server, is a potentially valid but less ideal choice compared to S3 or uploading. While HTTPS ensures secure transport, managing a separate web server adds complexity. S3 is generally the more straightforward AWS-native solution. Furthermore, direct upload simplifies the process and prevents dependency on an external server.

Option D, writing the script inline, is usually supported for smaller scripts. Inline embedding is possible if the script's length meets the constraints of the Launch Wizard.

In conclusion, using an S3 URL (C) leverages AWS's object storage service for secure and scalable script delivery, while direct upload (E) offers simplicity and eliminates external dependencies.

Question: 28

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:

The correct answer is ADE. Here's why:

A. Download the SAP HANA software. Launch Wizard requires the SAP HANA software to be available for deployment. The SAP HANA software itself is proprietary and must be obtained directly from SAP, not from AWS. Launch Wizard facilitates deployment, it doesn't provide the software.

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. Launch Wizard uses the software packages that are stored in an Amazon S3 bucket. It is vital to put the SAP HANA software in a accessible S3 bucket following AWS Launch Wizard's recommended file structure for correct software extraction and installation. Using an S3 bucket provides scalable and reliable storage for the software during deployment.

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. Launch Wizard requires an operating system AMI to build the instances. Either creating a custom AMI, or use an approved AMI for SAP from the AWS Marketplace. BYOI allows for greater customization, while AWS Marketplace provides pre-configured images with required SAP certifications. This allows flexibility in choosing the base OS for SAP HANA.

Why other options are incorrect:

B. Download the AWS CloudFormation template for the SAP HANA deployment. Launch Wizard generates the CloudFormation template based on user inputs, meaning the SAP solutions architect doesn't need to download it manually. CloudFormation templates are created automatically by Launch Wizard based on selected options.

C. Download and extract the SAP HANA software. Upload the SAP HANA software to an FTP server that Launch Wizard can access. Launch Wizard accesses the SAP HANA software directly from an S3 bucket and handles the extraction process itself. Furthermore, uploading to an FTP server would be less efficient and less secure than using S3.

F. Create the SAP file system by using Amazon Elastic Block Store (Amazon EBS) before the deployment. Launch Wizard automates the file system creation on EBS volumes. The SAP solution architect doesn't need to create the file systems manually beforehand. Launch Wizard manages the creation and configuration of the file systems required for SAP HANA automatically, simplifying the deployment process.

Supporting Documentation:

AWS Launch Wizard for SAP

Question: 29

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 VirtuallPAddress 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:

Here's a detailed justification for why options B and D are the correct steps for the SAP solutions architect to take, when deploying SAP HANA on AWS with Multi-AZ using AWS Launch Wizard and SUSE Linux Enterprise High Availability Extension:

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.

The PaceMakerTag is used by the SUSE Linux Enterprise High Availability Extension (SLE HA) to identify the EC2 instances that are part of the Pacemaker cluster for SAP HANA high availability. If the same PaceMakerTag is used by other EC2 instances in the same AWS Region, it will lead to conflicts within the Pacemaker cluster configuration. Pacemaker might attempt to manage those unrelated instances, causing unpredictable behavior and failure of the HA setup. Ensuring uniqueness prevents such conflicts.

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.

The VirtualIPAddress is crucial for the SAP HANA high availability setup. When a failover occurs, this IP address is moved from the primary HANA instance to the secondary HANA instance, ensuring continuous access to the HANA database. If the VirtualIPAddress is within the VPC CIDR or conflicts with any existing IP addresses in the route table associated with the subnets, routing conflicts will prevent successful failover.

The application layer that is connecting to the HANA database will fail to connect if there is an IP address conflict, so make sure to use an IP address outside of the VPC and that is not in any other route tables.

Why other options are incorrect:

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. Launch Wizard for SAP handles the creation and configuration of EBS volumes. Defining volume types (gp3, io1, etc.) is part of the Launch Wizard's configuration process, not something you pre-create.

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. This is partially correct, but the relevant consideration is the same AWS account. While cross-account conflicts could be theoretically possible, they are less likely and not the primary concern. The vital part is the virtual hostname is not used in the same account because it can lead to network configuration and resolution errors.

E. Before starting the Launch Wizard deployment, set up the SUSE Linux Enterprise High Availability Extension network configuration and security group. The Launch Wizard automates the deployment and configuration of the SUSE Linux Enterprise High Availability Extension along with the required network configurations and security groups. Manually setting up these components beforehand defeats the purpose of using Launch Wizard.

Authoritative Links:

AWS Launch Wizard for SAP: General information on Launch Wizard.
SUSE Linux Enterprise High Availability Extension: Documentation on SLE HA.

Question: 30

A company that has SAP workloads on premises plans to migrate an SAP environment to AWS. The company is new to AWS and has no prior setup. The company has the following requirements:
The application server and database server must be placed in isolated network configurations.
SAP systems must be accessible to the on-premises end users over the internet.
The cost of communications between the application server and the database server must be minimized. Which combination of steps should an SAP solutions architect take to meet these requirements? (Choose two.)

A.Configure a Network Load Balancer for incoming connections from end users.

B.Set up an AWS Site-to-Site VPN connection between the company's on-premises network and AWS.

C.Separate the application server and the database server by using different VPCs.

D.Separate the application server and the database server by using different subnets and network security groups within the same VPC.

E.Set up an AWS Direct Connect connection with a private VIF between the company's on-premises network and AWS.

Answer: BD

Explanation:

The correct answer is **BD**. Here's a detailed justification:

B: Set up an AWS Site-to-Site VPN connection between the company's on-premises network and AWS. This fulfills the requirement that SAP systems must be accessible to on-premises end users over the internet in a secure and cost-effective manner. A Site-to-Site VPN creates an encrypted tunnel, allowing secure communication between the on-premises network and the AWS environment. While Direct Connect (option E) offers lower latency and higher bandwidth, it's generally more expensive and might be overkill for a company new to AWS. A Site-to-Site VPN provides a good balance between security, cost, and ease of implementation for initial migration.

AWS Site-to-Site VPN

D: Separate the application server and the database server by using different subnets and network security groups within the same VPC. This ensures network isolation within AWS while minimizing communication costs. Placing the application and database servers in different subnets allows for the application of distinct network security groups. Network security groups act as virtual firewalls, controlling inbound and outbound traffic at the instance level. Using different security groups allows for restricting access to the database server only to the application server, enhancing security. Because the resources are within the same VPC, they can communicate using private IP addresses, which incurs no extra cost. Separating the application server and the database server into different VPCs (option C) would increase the cost and complexity due to the need for VPC peering or Transit Gateway, thus violating the minimizing cost constraint.

VPC Security Groups
Subnets

Option A is incorrect because it does not provide secure access to the SAP systems from the on-premise network. A Network Load Balancer is primarily used for distributing incoming application traffic across multiple targets. It might be used in conjunction with a VPN, but it doesn't replace the VPN itself. Option E (AWS Direct Connect) could be a better choice in the long run, it's a costlier alternative to Site-to-Site VPN and less appropriate for a company new to AWS.

In summary, setting up a Site-to-Site VPN provides secure access for on-premises users, and using different subnets and security groups within the same VPC allows for network isolation while minimizing communication costs within the AWS environment, fulfilling all the requirements.

Question: 31

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:

Here's a detailed justification for why option D is the best solution:

The key to answering this question lies in understanding which AWS services provide the right combination of functionality to address all the security requirements specified.

AWS Config is designed for evaluating, auditing, and assessing the configurations of your AWS resources. It enables you to establish and maintain a baseline configuration for your resources and continually monitor compliance against that baseline. AWS Config comes with **managed rules** for common configuration checks (e.g., checking for encrypted EBS volumes or specific instance types). However, it also allows you to create **custom rules** for more specialized checks using AWS Lambda functions, which is crucial for the SAP-specific parameter validation requirement.

EventBridge (formerly CloudWatch Events) provides a serverless event bus that allows you to react to changes in your AWS environment. You can configure EventBridge to trigger actions (like sending notifications) based on specific events, such as AWS Config flagging a resource as non-compliant. **Amazon SNS** is a notification service that works seamlessly with EventBridge to send out emails or other types of alerts.

Option A is not ideal because AWS AppConfig primarily focuses on managing application configuration rather than performing comprehensive security and compliance checks on infrastructure resources. While it can model configuration data, it lacks the specialized compliance evaluation capabilities of AWS Config. Also, using Systems Manager Automation for constant compliance monitoring might be overly complex and less

efficient than AWS Config's continuous evaluation.

Option B is good, but doesn't address the specialized SAP configuration checks. AWS Config managed rules are not going to be able to check SAP system parameters.

Option C is incorrect because AWS Trusted Advisor offers a limited set of high-level checks and does not cover the detailed, continuous compliance monitoring required by the scenario, particularly the custom rules needed for SAP system parameters. It's more for best practice recommendations, rather than active monitoring.

Option D, by combining AWS Config's managed rules for general AWS resource compliance with custom rules for SAP parameters and utilizing EventBridge and SNS for notifications, provides the most comprehensive and automated solution.

Supporting Links:

AWS Config:https://aws.amazon.com/config/
AWS Config Rules:https://docs.aws.amazon.com/config/latest/developerguide/evaluate-config.html
Amazon EventBridge:https://aws.amazon.com/eventbridge/
Amazon SNS:https://aws.amazon.com/sns/

Question: 32

A company is hosting its SAP workloads on AWS. An SAP solutions architect is designing high availability architecture for the company's production SAP S/4HANA and SAP BW/4HANA workloads. These workloads have the following requirements:
Redundant SAP application servers that consist of a primary application server (PAS) and an additional application server (AAS)
ASCS and ERS instances that use a failover cluster
Database high availability with a primary DB instance and a secondary DB instance
How should the SAP solutions architect design the architecture to meet these requirements?

A.Deploy ASCS and ERS cluster nodes in different subnets within the same Availability Zone. Deploy the PAS instance and AAS instance in different subnets within the same Availability Zone. Deploy the primary DB instance and secondary DB instance in different subnets within the same Availability Zone. Deploy all the components in the same VPC.

B.Deploy ASCS and ERS cluster nodes in different subnets within the same Availability Zone. Deploy the PAS instance and AAS instance in different subnets within the same Availability Zone. Deploy the primary DB instance and secondary DB instance in different subnets within the same Availability Zone. Deploy the ASCS instance, PAS instance, and primary DB instance in one VPC. Deploy the ERS instance, AAS instance, and secondary DB instance in a different VPC.

C.Deploy ASCS and ERS cluster nodes in different subnets across two Availability Zones. Deploy the PAS instance and AAS instance in different subnets across two Availability Zones. Deploy the primary DB instance and secondary DB instance in different subnets across two Availability Zones. Deploy all the components in the same VPC.

D.Deploy ASCS and ERS cluster nodes in different subnets across two Availability Zones. Deploy the PAS instance and AAS instance in different subnets across two Availability Zones. Deploy the primary DB instance and secondary DB instance in different subnets across two Availability Zones. Deploy the ASCS instance, PAS instance, and primary DB instance in one VPC. Deploy the ERS instance, AAS instance, and secondary DB instance in a different VPC.

Answer: C

Explanation:

The correct answer is C. Here's why:

High availability in AWS hinges on distributing components across multiple Availability Zones (AZs). An AZ is

a physically isolated location within an AWS Region. Deploying across AZs ensures that if one AZ fails, the application can continue to operate in another AZ.

Option C deploys ASCS and ERS, PAS and AAS, and the primary and secondary database instances across two AZs. This aligns with best practices for high availability in AWS. Keeping all the components within the same VPC simplifies networking and communication between the various SAP components.

Options A and B place all or some components within the same Availability Zone, negating the benefit of redundancy in case of an AZ failure. Option D divides the components across different VPCs. While technically possible, this complicates the network architecture significantly and adds unnecessary complexity, making it more difficult to manage the SAP landscape. Furthermore, communication between components is more complex across VPCs compared to within a single VPC.

Therefore, distributing all components across different AZs within a single VPC as in option C, provides the best balance of high availability, fault tolerance, and simplified management.

Supporting resources:

AWS Regions and Availability Zones: Explains the concept of AZs and Regions
SAP on AWS: Outlines best practices for deploying SAP on AWS
AWS VPC: Describes how to use VPCs for network isolation. While VPC peering is possible it adds complexity, and isn't required for this scenario.

Question: 33

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-1:11111111111: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:

B. Change the Action element to ec2:ReplaceRoute.

Reference:

https://docs.aws.amazon.com/sap/latest/sap-hana/sap-hana-on-aws-oip.html

Question: 34

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).

Answer: D

Explanation:

Here's a detailed justification for why option D is the most cost-effective solution for the SAP disaster recovery (DR) scenario, focusing on meeting the RPO/RTO objectives while considering cost and existing infrastructure:

SAP HANA Database Replication (HSR)

Lower Cost: SAP HANA System Replication (HSR) is generally a more cost-effective DR solution compared to solutions like CloudEndure for the database layer because it leverages the built-in capabilities of SAP HANA itself.
RPO Compliance: Asynchronous HSR provides an RPO close to zero, which easily satisfies the 15-minute RPO requirement. HSR's asynchronous mode ensures minimal performance impact on the primary production system.
EC2 Instance Sizing: Using a smaller SAP-certified EC2 instance for DR is cost-effective. This DR instance doesn't need to match the production appliance's performance until a failover event occurs. You can scale up the EC2 instance size during a DR event.

No Data/Log Backup Changes: The solution avoids changes to the existing on-premises data and log backup configurations, minimizing disruption.

SAP Application Server DR with CloudEndure

Efficient Replication: CloudEndure Disaster Recovery provides block-level, continuous data replication of the SAP application servers to AWS.

Fast RTO: CloudEndure's low-cost, continuous replication achieves minimal RTO since the DR servers are constantly kept up-to-date.

VM Migration Alternative: CloudEndure offers a more robust and manageable solution compared to one-time VM Import/Export, which can become cumbersome for continuous DR.

NFS File Share Replication

AWS DataSync and Amazon EFS: AWS DataSync is purpose-built for efficiently and securely transferring data over the network to AWS storage services like Amazon EFS.

Real-Time Synchronization: DataSync enables real-time or near real-time synchronization of /sapmnt and /usr/sap/trans to Amazon EFS, which ensures that changes to the transport directory are available in the DR environment.

Why Other Options are Less Suitable

Option A & B (S3 & Storage Gateway): Relying on S3 and AWS Storage Gateway for HANA data and log backups is feasible but less efficient for DR scenarios, especially given the 15-minute RPO requirement. The recovery process from S3 can be lengthy and complex, potentially exceeding the 4-hour RTO. It also ignores the efficiency of HSR.

Option C (CloudEndure for everything): This is a functional but less cost-effective solution. It's overkill to replicate the database with CloudEndure when SAP HANA's built-in HSR offers a more efficient and cheaper alternative.

Authoritative Links:

SAP HANA System Replication:
https://help.sap.com/docs/HANA_PLATFORM/6b94445c962c4e78a1b3f9579eadb9fe/20a8d6a475191014b0a286 **AWS DataSync:**https://aws.amazon.com/datasync/
CloudEndure Disaster Recovery:https://aws.amazon.com/cloudendure-disaster-recovery/

Question: 35

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:

The correct answer is D, Amazon FSx for Windows File Server. Here's why:

The scenario involves migrating a Windows-based SAP environment with file shares used for transport directories and third-party application data. These requirements immediately point towards a Windows-compatible file system service.

Amazon FSx for Windows File Server: This service provides fully managed, native Microsoft Windows file systems built on Windows Server. It supports SMB protocol, Active Directory integration, and NTFS permissions, which are critical for compatibility with existing Windows-based SAP applications and file share infrastructure. This ensures a seamless "lift and shift" migration while adhering to AWS best practices.

Let's analyze why the other options are not ideal:

Amazon EBS: EBS volumes are block storage intended for persistent storage attached to a single EC2 instance. While you could create a file share on an EBS volume, it's not designed for shared access from multiple instances, especially within an SAP environment. This solution is complex to manage and lacks the features of a dedicated file service.

AWS Storage Gateway: Storage Gateway connects on-premises applications to AWS cloud storage. While useful for hybrid scenarios and data backup, it doesn't directly host the file shares in the AWS cloud. The company's goal is to migrate entirely to AWS, so a hybrid approach is not needed in this specific context.

Amazon EFS: EFS provides a scalable, elastic, cloud-native NFS file system. While excellent for Linux-based workloads, it doesn't natively support the SMB protocol or Windows-specific features needed for SAP on Windows. It is not designed to integrate with Windows Active Directory.

FSx for Windows File Server, in contrast, handles these issues elegantly, facilitating integration with Active Directory for user authentication and authorization. It also provides features like data deduplication to save on storage costs. The ability to manage the file server within the AWS Management Console and integration with other AWS services make it the ideal option in accordance with AWS best practices.

Authoritative links:

Amazon FSx for Windows File Server: https://aws.amazon.com/fsx/windows/
SAP on AWS Best Practices: https://aws.amazon.com/sap/

Question: 36

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: A

Explanation:

Here's a detailed justification for why option A is the best solution, along with supporting concepts and links:

Option A is the most cost-optimized and effective solution for meeting the company's DR requirements. It leverages SAP HANA System Replication (HSR) with preload off to maintain a near real-time copy of the SAP HANA database in the us-west-1 Region. The key advantage is the dual utilization of the secondary instance.

Using the secondary instance for QA during normal operations minimizes idle resources and reduces costs, addressing the "cost-optimized" requirement.

HSR ensures a Recovery Point Objective (RPO) of under 30 minutes, as data is constantly replicated. The "preload off" configuration allows for immediate data availability on the secondary instance without consuming unnecessary resources for production processing.

The Recovery Time Objective (RTO) is also met. The process of shutting down QA and starting production services during a DR event will be faster compared to launching new instances from AMIs or using a DR tool that requires conversion or staging. HSR is designed for quick failover.

Option B, while potentially viable, uses CloudEndure, which adds complexity and potentially higher costs compared to native HANA System Replication. The cutover and instance resizing steps introduce delays, potentially increasing RTO.

Option C, using HSR with preload on, would keep the secondary instance fully active as a hot standby. While it provides the fastest failover, it's the most expensive because it requires double the resources constantly running. This goes against the "cost-optimized" requirement.

Option D is the least suitable. Relying on AMIs, S3 replication, and point-in-time recovery introduces significant delays, making it difficult to achieve an RTO of less than 30 minutes. Manually updating host information and performing a point-in-time recovery are also time-consuming processes. Furthermore, this option involves more manual steps and is less automated compared to HSR. The RPO would also be significantly higher than 30 minutes as it relies on the frequency of backups.

In summary, option A provides the optimal balance of cost, RPO, and RTO by utilizing SAP HANA System Replication effectively and dual-purposing the secondary instance.

Supporting Links:

SAP HANA System Replication:
https://help.sap.com/docs/HANA_PLATFORM/6b94445c9626475c94eb9da33c65269c/f1e9c0389a564842ad03e **AWS and SAP:**https://aws.amazon.com/sap/

Question: 37

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 QAT 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 (io2) Amazon Elastic Block Store (Amazon EBS) volumes.
The entire development team 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?

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 (st1) for the /hana/data and /hana/log file systems for the production and non-production SAP HANA databases.

Answer: B

Explanation:

The correct answer is **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.**

Here's why:

Cost Optimization: The primary goal is to reduce infrastructure costs. Stopping non-production EC2 instances during off-peak hours (nights and weekends) is a highly effective cost-saving strategy. These instances aren't needed outside of business hours, so shutting them down eliminates the EC2 usage charges.

Least Administrative Effort: Amazon EventBridge allows scheduling events (like stopping and starting EC2 instances) automatically. This eliminates the need for manual intervention, minimizing administrative overhead. No manual scripts need to be maintained.

Suitable Instance Types: The development and sandbox environments are running on On-Demand Instances, which are ideal for variable workloads and situations where instances are not constantly running. This makes them suitable for stopping and starting.

Considerations for Reserved Instances: The SPT and production systems are running on Reserved Instances. Reserved Instances are cost-effective when you need consistent compute capacity for an extended period.

Since the production environment is business-critical and requires constant uptime, reserved instances are the perfect way to accomplish this. Because you are already paying for them, this is the proper place for them.

Why other options are not optimal:

A. Use a Spot Fleet instead of the Reserved Instances and On-Demand Instances: Spot Instances can be significantly cheaper, but they are subject to interruption if the Spot price exceeds your bid. This is unsuitable for production and performance testing systems, which need consistent availability. Production systems should use Reserved Instances, Dedicated Hosts, or Savings Plans for cost efficiency.

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: Requires constant involvement from staff, making it extremely inefficient and costly. EventBridge automates this process much more effectively.

D. Change the EBS volume type to Throughput Optimized HDD (st1) for the /hana/data and /hana/log file systems for the production and non-production SAP HANA databases: The EBS io2 type provides high IOPS, which is essential for SAP HANA. ST1 volumes are designed for less frequently accessed data. Using ST1 volumes for /hana/data and /hana/log would severely impact the performance of the SAP HANA systems.

References:

Amazon EventBridge: https://aws.amazon.com/eventbridge/
Amazon EC2 Instance Types: https://aws.amazon.com/ec2/instance-types/
Amazon EBS Volume Types: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volume-types.html

Question: 38

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
```

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:

Add the SAP HANA system to SAP HANA Studio. Select multiple container mode, and then try to initiate the backup again.

Reference:

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

Question: 39

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:

The correct answer is D: Four Reserved Instances and two On-Demand Instances. Here's why:

Reserved Instances (RIs) offer significant cost savings (up to 75% compared to On-Demand) in exchange for a commitment to use a specific instance type for a specified period (1 or 3 years). Since the company requires four application servers for the majority of the month during normal operations, purchasing four RIs is the most cost-effective approach for this baseline capacity.

The remaining two servers are only required during the monthly payroll and credit data processing peaks. On-Demand Instances are ideal for workloads that are short-term, spiky, or unpredictable. In this scenario, the company knows exactly when these peak workloads will occur and their duration is likely limited. Therefore, utilizing On-Demand Instances for these specific periods allows the company to scale up its compute capacity without incurring the long-term costs associated with Reserved Instances.

Option A is less suitable because Spot Instances can be terminated by AWS with little notice, which would risk interrupting the critical payroll and credit data processing tasks during peak times. Spot Instances are ideal for fault-tolerant or flexible workloads, which the peak SAP ECC tasks are not.

Option B is less cost-effective because On-Demand Instances are the most expensive purchasing option for consistent usage, even for the base load of four servers.

Option C is also less cost-effective because purchasing six RIs means paying for the capacity of two servers even during the periods of the month when they aren't needed.

Using a combination of RIs for the baseline capacity and On-Demand Instances for peak periods optimizes cost efficiency while guaranteeing availability for the critical SAP ECC workloads.

Relevant links for further research:

AWS EC2 Pricing:https://aws.amazon.com/ec2/pricing/
AWS Reserved Instances:https://aws.amazon.com/ec2/reserved-instances/ **AWS**
On-Demand Instances:https://aws.amazon.com/ec2/pricing/on-demand/ **AWS**
Spot Instances:https://aws.amazon.com/ec2/spot/