Practice Exam Questions

amazon web services

PAS-C01

AWS Certified SAP on AWS - Specialty

EXAMKILLER

Help Pass Your Exam At First Try

Amazon Web Services

Exam PAS-C01

AWS Certified: SAP on AWS - Specialty

Version: 3.2

[Total Questions: 65]

Question No : 1

A company needs to implement high availability for its SAP S 4HANA system on AWS The company will use a SUSE Linux Enterprise Server clustering solution in private subnets across two Availability Zones An SAP solutions architect must ensure that the solution can route traffic to the active SAP instance m this clustered configuration.

What should the SAP solutions architect do to meet these requirements?

A. Implement the SAP cluster solution by using a secondary private IP address Reassign the secondary private IP address from one network interface to another network interface in the event of any failure that affects the primary instance.
B. Implement the SAP cluster solution by using an Elastic IP address Mask the failure of an instance or software by rapidly remapping the address to another instance in the account
C. Implement the SAP duster solution by using a public IP address Use this public IP address for communication between the instances and the internet
D. implement the SAP duster solution by using an overlay IP address that is outside the CIDR block of the VPC Use overlay IP address routing to dynamically update the route table to point to the active node and provide external access by using a Network Load Balancer or AWS Transit Gateway.

Answer: D

Question No : 2

A company is running an SAP ERP Central Component (SAP ECC) system on an SAP HANA database that is 10 TB m size The company rs 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 mat 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 UploadChanneiSize 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 tor 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 every day, it is likely that the backups are taking too long to complete because the upload channel size is not sufficient for the size of the backups. By increasing the UploadChannelSize parameter, the SAP NetWeaver administrator can adjust the amount of data that is sent over the network at a time, which can help to speed up the backups and reduce the time they take to complete. This can be done by editing the aws-backint-agent-config yaml configuration file and increasing the value of the UploadChannelSize parameter based on the EC2 instance type and storage configurations.

Question No : 3

A company runs its SAP ERP 6 0 EHP 8 system on SAP HANAon AWS The system is deployed on an r4 I6xlarge Amazon EC2 instance with default tenancy. The company needs to migrate the SAP HANA database to an x2gd/.6xiarge High Memory instance After an operations engineer changes the instance type and starts the instance the AWS Management Console shows a failed instance status check

What is the cause of this problem?

A. The operations engineer missed the network configuration step during the post-migration activities
B. The operations engineer missed the Amazon CloudWatch configuration step during the post-migration activities.
C. The operations engineer did not install Elastic Network Adapter (ENA) drivers before changing the instance type
D. The operations engineer did not create a new AMI from the original instance and did not launch a new instance with dedicated tenancy from the AMI

Answer: C

Explanation: The Elastic Network Adapter (ENA) is a software-based network interface that provides high-performance network connectivity and is required for instances with higher network performance requirements. If the ENA drivers are not installed before changing the instance type, the instance will not be able to communicate with the network, resulting in a failed instance status check.

Question No : 4

A company hosts its SAP NetWeaver workload on SAP HANA m the AWS Cloud The SAP NetWeaver application is protected by a cluster solution that uses Red Hat Enterprise Linux High Availability Add-On The duster solution uses an overlay IP address to ensure that the high availability cluster is still accessible during failover scenarios.

An SAP solutions architect needs to facilitate the network connection to this overlay IP address from multiple locations These locations include more than 25 VPCs other AWS Regions and the on-premises environment The company already has set up an AWS Direct Connect connection between the on-premises environment and AWS.

What should the SAP solutions architect do to meet these requirements in the MOST scalable manner?

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

Answer: B

Explanation: AWS Transit Gateway allows the SAP solutions architect to connect multiple VPCs and on-premises networks together in a scalable manner. It acts as a hub that controls how traffic is routed between the connected networks. By attaching the VPCs and the on-premises environment to the Transit Gateway, the SAP solutions architect can establish a single connection to the overlay IP address in the high availability cluster, ensuring that the cluster is accessible from all locations.

Question No : 5

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

When solution will meet these requirements?

A. Deploy a single EC2 instance to support the secondary database in us-west with additional storage Use this secondary database instance to support QA and production Configure the primary SAP HANA database in us-east-1 to constantly replicate the data to the secondary SAP HANA database in us-west-t by using SAP HANA system replication with preload off During DR shut down the QA SAP HANA instance and restart the production services at the secondary site
B. Deploy a secondary staging server on an EC2 instance in us-west-1 Use CloudEndure Disaster Recovery to replicate changes at the database level from us-east-1 to the

secondary staging server on an ongoing basis During DR, initiate cutover increase the size of the secondary EC2 instance to match the primary EC2 instance and start the secondary EC2 instance

C. Set up the primary SAP HANA database in us-east-1 to constantly replicate the data to a secondary SAP HANA database in us-west-1 by using SAP HANA system replication with preload on Keep the secondary SAP HANA instance as a hot standby that 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 m us-west-1 During DR launch the EC2 instance in us-west-1 based on AMIs that are replicated Update host information Download database and log backups from the secondary S3 bucket Perform a point-in-time recovery

Answer: C

Explanation: 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. This solution will allow the company to have a secondary HANA instance that is always ready to take over in case of failure, with an RPO and RTO of less than 30 minutes. The solution also allows the company to have a guaranteed capacity reservation and cost-optimized solution.

Question No : 6

An SAP technology consultant needs to scale up a primary application server (PAS) instance The PAS currently runs on a c5a.xlarge Amazon EC2 instance The SAP technology consultant needs to change the instance type to c5a 2xlarge

How can the SAP technology consultant meet this requirement?

A. Stop the complete SAP system Stop the EC2 instance Use the AWS Management Console or the AWS CLI to change the instance type Start the EC2 instance Start the complete SAP system

B. While SAP is running use the AWS Management Console or the AWS CLI to change the instance type without stopping the EC2 instance

C. Stop the complete SAP system Terminate the EC2 instance Use the AWS Management Console or the AWS CLI to change the instance type Start the EC2 instance Start the complete SAP system

D. While SAP is running log in to the EC2 instance. Run the following AWS CLI command:
aws ec2 modify-instance-attribute –instance-id <INSTANCEID>
--instance-type “{\”Value\”: |”c5a.2xlarge1\”}”.

Answer: B

Question No : 7

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 win 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 (1AM) to deny access from the IP address block
D. Configure the firewall m the operating system of the EC2 instances to deny access from the IP address block

Answer: A

Explanation: The company can meet its requirement by modifying the network access control lists (ACLs) that are associated with all public subnets in the VPC to deny access from the offending IP address block. This would deny access to the VPC from all the IP addresses that are attempting port scans, and would be effective for the next 24 hours. Security groups are associated with individual instances, it would be more time-consuming to update all instances security groups and it's not scalable. AWS Identity and Access Management (IAM) is mainly used to manage user access to AWS resources and it's not appropriate for this use case. Configuring the firewall on the operating system of the EC2 instances would be less effective as it does not provide a centralized and scalable solution for managing access control across all subnets in the VPC.
Top of Form

Question No : 8

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.

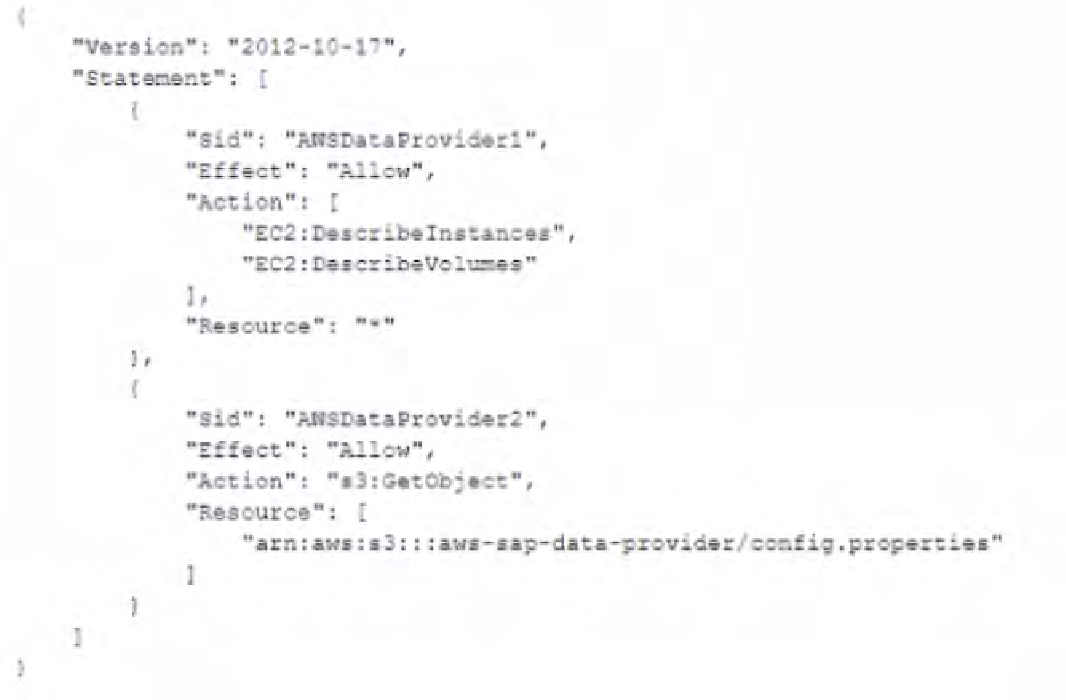

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AWSDataProvider1",
            "Effect": "Allow",
            "Action": [
                "EC2:DescribeInstances",
                "EC2:DescribeVolumes"
            ],
            "Resource": "*"
        },
        {
            "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 1AM permissions to resolve this issued.

A. Add the cloudwatch ListMetrics action to the policy statement with Sid AWSDataProvider1.
B. Add the cloudwatch GetMetricStatrstics action to the policy statement with Sid AWSDataProvider1
C. Add the cloudwatch GetMetricStream action (o the policy statement with Sid AWSDataProvider
D. Add the cloudwatch DescribeAlarmsForMetric action to the policy statement with Sid AWSDataProvider

Answer: B

Explanation: The AWS Data Provider for SAP requires the ability to access metrics data in order to return metrics to the SAP application. The IAM policy statement with Sid "AWSDataProvider1" currently does not have the necessary permissions to access metrics data. The SAP solutions architect should add the cloudwatch:GetMetricStatistics action to the policy statement with Sid "AWSDataProvider1" to grant the necessary permissions for the Data Provider to access metrics data.
The other actions such as "EC2:DescribeInstances" and "EC2:DescribeVolumes" are not related to CloudWatch metrics and only provide the ability to describe EC2 instances and volumes. Actions such as "s3:GetObject" are not related to CloudWatch metrics, it's used to get an object from an S3 bucket. Actions such as "cloudwatch:ListMetrics" and "cloudwatch:DescribeAlarmsForMetric" would not be necessary for the AWS Data Provider for SAP to return metrics to the SAP application and it's not related to the problem described.

Question No : 9

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 FODN 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 VPC Add 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

Question No : 10

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

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

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

D. Stop the EC2 instances

Answer: A

Explanation: Ensure that the SAP kernel versions are 7.45 or later. This is important because IMDSv2 is only supported by SAP kernel versions 7.45 and later. If the SAP kernel versions are not at least 7.45, then the enhancement will cause an SAP outage as the instances will not be able to communicate with the metadata service.

Question No : 11

A company is hosting its SAP workloads on AWS An SAP solutions architect is designing high availability architecture for the company's production SAP S4HANA 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 OB 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 m different subnets within the same Availability Zone Deploy all the components in the same VPC
B. Deploy ASCS and ERS duster nodes m different subnets within the same Availability Zone Deploy the PAS instance and AAS instance in different subnets within the same Availably Zone Deploy the primary DB instance and secondary DB instance m 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 m 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 m 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: This solution would ensure that the ASCS and ERS instances are deployed in different subnets across different Availability Zones, providing redundancy for the failover cluster. The PAS and AAS instances are also deployed in different subnets across different Availability Zones, providing redundancy for the application servers. The primary and secondary DB instances are also deployed in different subnets across different Availability Zones, providing redundancy for the database. Additionally, all the components are deployed in the same VPC, which will minimize the cost of communication between the application server and the database server.

Question No : 12

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 dies sapmnt and . usrsap. trans are hosted on an Amazon Elastic File System (Amazon EFS) tile 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? (Select 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: A,D

Explanation: The Application Load Balancer (ALB) would be a good solution for the company's requirements, it provides a layer-7 load balancing and it's highly available, it allows the company to use the overlay IP address as a target and makes the SAP NetWeaver system accessible through the SAP GUI client tool.
AWS Transit Gateway can also be used to meet the company's requirements, it provides a centralized and scalable solution to route traffic between VPCs. The company can configure the overlay IP address as a static route in the transit gateway route table and specify the VPC as a target. This would allow the company to use the existing private connectivity to the VPC and make the SAP NetWeaver system accessible through the SAP

GUI client tool.

Question No : 13

An SAP engineer has deployed an SAP S 4HANA system on an Amazon EC2 instance mat runs Linux. The SAP license key has been installed After a white the newly installed SAP instance presents an error that indicates that the SAP license key is not valid because the SAP system's hardware key changed. There have been no changes to the EC2 instance or its configuration.

Which solution will permanently resolve this issue?

A. Perform SAP kernel patching
B. Apply a new SAP license that uses a new hardware key Install the new key
C. Set the SUC_HW_VERSlON Linux environment variable
D. Reboot the EC2 instance

Answer: B

Explanation: When the hardware key changes, the license key becomes invalid and it requires to install new license key that uses the new hardware key.

Question No : 14

A company is implementing SAP HANA on AWS According 10 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: This is the most operationally efficient solution that meets the company's security policy requirements. AWS KMS is a service that enables you to create and manage encryption keys that are used to encrypt and decrypt data. By configuring AWS Backup Agent for SAP HANA to use AWS KMS for SAP backups, the company can ensure that the backups are encrypted at rest and that only authorized team members have the ability to decrypt them. The key policy allows the company to define which team members are authorized to access the key, so that it can be used to decrypt the backup. This approach is operationally efficient because it does not require the company to manually encrypt and decrypt backups, and it enables the company to manage access to the encryption key through IAM policies, without the need for sharing encryption keys.

Question No : 15

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

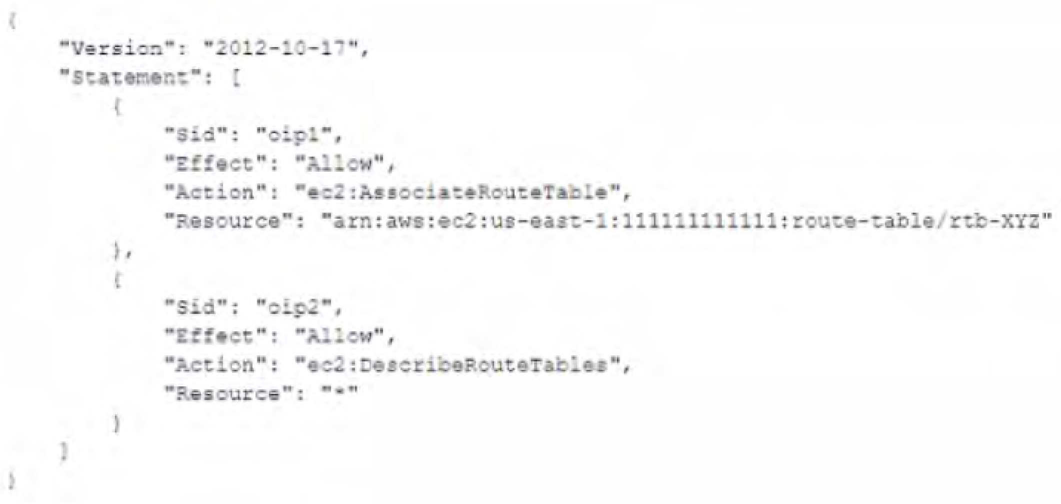

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "oip1",
            "Effect": "Allow",
            "Action": "ec2:AssociateRouteTable",
            "Resource": "arn:aws:ec2:us-east-1:111111111111:route-table/rtb-XYZ"
        },
        {
            "Sid": "oip2",
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

During a test of failover the SAP solutions architect finds that the overlay IP address does not change to the secondary Availability Zone Which change should the SAP solutions architect make in the policy statement for Sid oip1 to dx this error1?

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

D. Change the Action element to ec2 ReplaceTransrtGatewayRoute

Answer: B

Question No : 16

A company uses an SAP application that runs batch jobs that ate performance sensitive. The batch jobs can be restarted safely The SAP application has sot application servers. The SAP application functions reliability as long as the SAP application availability remains greater than 60%. The company wants to migrate the SAP application to AWS. The company is using a duster with two Availability Zones

How should the company distribute the SAP application servers to maintain system reliability?

A. Distribute the SAP application servers equally across three partition placement groups
B. Distribute the SAP application servers equally across three Availability Zones
C. Distribute the SAP application servers equally across two Availability Zones
D. Create an Amazon EC2 Auto Scaling group across two Availability Zones Set a minimum capacity value of 4.

Answer: B

Explanation: This will ensure that even in the event of a failure in one Availability Zone, the remaining two Availability Zones will still have sufficient capacity to meet the availability requirement of greater than 60%. This will also provide an additional layer of redundancy and protection against a single point of failure.

Question No : 17

A company wants 10 run SAP HANA on AWS m the eu-centrai-1 Region. The company must make the SAP HANA system highly available by using SAP HANA system replication in addition the company must create a disaster recovery (DR) solution that uses SAP HANA system replication in the eu-west-1 Region As prerequisites the company has confirmed that inter-AZ latency is less than 1 ms and that Inter-Region latency is greater than 1 ms.

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

A. Install the tier 1 primary system and the tier 2 secondary system in eu-central-1 Configure the tier 1 system m Availability Zone 1 Configure the tier 2 system m Availability Zone 2 Configure SAP HANA system replication between tier 1 and tier 2 by using ASYNC replication mode install the OR tier 3 secondary system m 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 m Availability Zone 2 Configure SAP HANA system replication between tier 1 and tier 2 by using SYNC replication mode Install the OR her 3 secondary system n eu-west-1 by using ASYNC replication mode.

C. Install the tier 1 primary system and the tier 2 secondary system in eu-central-1 Configure the tier 1 system in Availability Zone 1 Configure the tier 2 system in Availability Zone 2 Configure SAP HANA system replication between tier 1 and tier 2 by using SYNC replication mode Install the OR tier 3 secondary system n eu-west-1 Store daily backups from tier 1 m an Amazon S3 bucket m eu-central-1 Use S3 Cross-Region Replication to copy the daily backups to eu-west-i 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 m 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 m Availability Zone 1 Configure the tier 2 system m 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: A,C

Explanation: Ensures high availability and disaster recovery by using SAP HANA system replication in two different availability zones in eu-central-1, and then installing a third secondary system in eu-west-1 with SYNC replication mode, which provides a fallback option in case of disaster, also it meets the inter-AZ latency requirement.

Ensures high availability and disaster recovery by using SAP HANA system replication in two different availability zones in eu-central-1, and then storing daily backups from tier 1 in an Amazon S3 bucket in eu-central-1 and then using S3 cross-region replication to copy the backups to eu-west-1 where they can be restored if needed, this meets the inter-region latency requirement.

Question No : 18

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