DUMPSARENA

AWS Certified: SAP on AWS - Specialty

Amazon AWS PAS-C01

Version Demo

Total Demo Questions: 10

Total Premium Questions: 65

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

A company has deployed its SAP applications into multiple Availability Zones in the same AWS Region To accommodate storage of media files database table export and import and files dropped by third-party tools the company has mounted Amazon Elastic File System (Amazon EFS) file systems between the SAP instances The company needs to retrieve the files quickly for installations updates and system refreshes

Over time the EFS file systems have grown exponentially to multiple terabytes An SAP solutions architect must optimize storage cost for the tiles that are stored in Amazon EFS.

Which solution will meet this requirement with the LEAST administrative overhead?

A. Scan the files manually to identify unnecessary files Delete the unnecessary files

B. Move the files to Amazon S3 Glacier Deep Archive

C. Apply a lifecycle policy on the files in Amazon EFS to move the files to EFS Standard-Infrequent Access (Standard-IA)

D. Move the files to Amazon S3 Glacier Apply an S3 Glacier vault lock policy to the files

ANSWER: C

Explanation:

Apply a lifecycle policy on the files in Amazon EFS to move the files to EFS Standard-Infrequent Access (Standard-IA) is the solution that will meet this requirement with the least administrative overhead. This solution allows you to automatically move files that have not been accessed for a certain period of time to a lower-cost storage class, reducing storage costs without the need for manual intervention.

QUESTION NO: 2

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: 3

A company is migrating its SAP workloads to AWS The company's IT team installs a highly available SAP S.4HANA system that uses the SAP HANA system replication cluster package on SUSE Linux Enterprise Server. The IT team deploys the system by using cluster nodes m different Availability Zones within the same AWS Region.

After the initial launch of the SAP application the application is accessible However after failover the IT team cannot access the application even though the system is up and running on the secondary node After investigation an SAP solutions architect discovers that the virtual IP address has not been used correctly

Which combination of steps should the SAP solutions architect take to resolve this problem? (Select TWO.)

A. Use an overlay IP address as a secondary IP address with the primary node of the cluster

B. Choose an overlay IP address within the VPC CiDR block that corresponds with the secondary node of the cluster

C. Use an overlay IP address as a virtual IP address

D. Choose an overlay IP address within the VPC CIDR block that corresponds with the primary node of the cluster

E. Choose an overlay IP address outside the VPC CiDR block that hosts the application and the database

ANSWER: A D

QUESTION NO: 4

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 tor 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: B

Explanation:

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 would be the most cost-effective DR option for the company's SAP HANA database. This would allow the company to save money on the DR process by using a less powerful instance type for the secondary instance and also replication with preload option turned off will reduce the cost of data transfer.

QUESTION NO: 5

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 an 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: A

Explanation:

This allow the company to have a guaranteed capacity for the four application servers that are needed during normal operation, and also utilize Spot Instances for the additional two servers during peak usage. Spot Instances can save up to 90% on the cost of On-Demand instances and thus it can save the company a lot of cost during the peak usage when it only needs two additional servers.

QUESTION NO: 6

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 poor 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? (Select 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: C E

Explanation:

To meet the requirements, an SAP solutions architect should take the following steps:

Using a Network Load Balancer for incoming connections from end users is not a necessary step, as it's not required to meet the requirements Setting up an AWS Site-to-Site VPN connection is not a necessary step, as it's not required to meet the requirements Separating the application server and the database server by using different subnets and network security groups within the same VPC is not a necessary step as it's not required to meet the requirements

QUESTION NO: 7

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 m 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 >n 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 9MB.

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 lo 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:

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 and Configure an S3 Lifecycle policy to send the backup files to S3 Glacier based on the company's data retention policy. This option is cost-effective because it avoids the need for third-party tools, tape drives and storage gateways, and reduces the amount of time and resources needed for the migration process. Additionally, the S3 lifecycle policy allows you to automate the storage and archiving process and ensure that your data is stored in the most cost-effective way.

QUESTION NO: 8

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: 9

A company wants 10 migrate its SAP ERP landscape to AWS The company will use a highly available distributed deployment for the new architecture Clients will access SAP systems from a local data center through an AWS Site-to-Site VPN connection that is already in place An SAP solutions architect needs to design the network access to the SAP production environment

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

A. For the ASCS instance configure an overlay IP address that is within the production VPC CIDR range Create an AWS Transit Gateway Attach me 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: A D

Explanation:

AWS Transit Gateway provides a central hub for the traffic between different VPCs, on-premises data centers, and AWS services. By attaching the VPN connection to the Transit Gateway, the solutions architect can route the communication between the local data center and the production VPC. Additionally, By configuring an overlay IP address within the production VPC, the solutions architect can create a static route on the production VPC to route the traffic to the ASCS instance.

An Application Load Balancer can also provide high availability to the ASCS instance by distributing the traffic to the instance. Additionally, by configuring an overlay IP address outside the production VPC, the solutions architect can create

QUESTION NO: 10

A company has moved an of its SAP workloads to AWS During peak business hours end users are reporting performance issues because work processes are going into PRIV mode on an SAP S/4HANA system An SAP support engineer indicates that SAP cannot provide support for this issue because some specific performance metrics are not available.

Which combination of actions must the company perform to comply with SAP support requirements? (Select THREE.)

A. Buy an SAP license from AWS Ensure that the SAP license is installed

B. Select only an AWS Migration Acceleration Program (MAP) certified managed service provider (MSP)

C. Enable detailed monitoring for Amazon CloudWatch on each Amazon EC2 instance where SAP workloads are running

D. Install configure and run the AWS Data Provider for SAP on each Amazon EC2 instance where SAP workloads are running

E. integrate AWS Systems Manager with SAP Solution Manager to provide alerts about SAP parameter configuration drift

F. Enable SAP enhanced monitoring through a SAPOSCOL enhanced function

ANSWER: A C D