Free Questions for SAP-C02 by dumpshq

Shared by Hicks on 24-05-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

A company is running an Apache Hadoop cluster on Amazon EC2 instances. The Hadoop cluster stores approximately 100 TB of data for weekly operational reports and allows occasional access for data scientists to retrieve dat

a. The company needs to reduce the cost and operational complexity for storing and serving this data.

Which solution meets these requirements in the MOST cost-effective manner?

Options:

A) Move the Hadoop cluster from EC2 instances to Amazon EMR. Allow data access patterns to remain the same.

B) Write a script that resizes the EC2 instances to a smaller instance type during downtime and resizes the instances to a larger instance type before the reports are created.

C) Move the data to Amazon S3 and use Amazon Athena to query the data for reports. Allow the data scientists to access the data directly in Amazon S3.

D) Migrate the data to Amazon DynamoDB and modify the reports to fetch data from DynamoDB. Allow the data scientists to access the data directly in DynamoDB.

Answer:

C

Explanation:

'The company needs to reduce the cost and operational complexity for storing and serving this data. Which solution meets these requirements in the MOST cost-effective manner?' EMR storage is ephemeral. The company has 100TB that need to persist, they would have to use EMRFS to backup to S3 anyway. https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-plan-storage.html

100TB

EBS - 8.109$

S3 - 2.355$

You have saved 5.752$

This amount can be used for Athen. BTW. we don't know indexes, amount of data that is scanned. What we know is that tit will be: 'occasional access for data scientists to retrieve data'

Question 2

Question Type: MultipleChoice

A company wants to retire its Oracle Solaris NFS storage arrays. The company requires rapid data migration over its internet network connection to a combination of destinations for Amazon S3. Amazon Elastic File System (Amazon EFS), and Amazon FSx lor Windows File Server. The company also requires a full initial copy, as well as incremental transfers of changes until the retirement of the storage arrays. All data must be encrypted and checked for integrity.

What should a solutions architect recommend to meet these requirements?

Options:

A) Configure CloudEndure. Create a project and deploy the CloudEndure agent and token to the storage array. Run the migration plan to start the transfer.

B) Configure AWS DataSync. Configure the DataSync agent and deploy it to the local network. Create a transfer task and start the transfer.

C) Configure the aws S3 sync command. Configure the AWS client on the client side with credentials. Run the sync command to start the transfer.

D) Configure AWS Transfer for FTP. Configure the FTP client with credentials. Script the client to connect and sync to start the transfer.

Answer:

B

Explanation:

It enables secure, high-performance transfers and supports both full initial copies and incremental transfers of changes. DataSync provides encryption and checksum validation to ensure data integrity, and it can be configured to transfer data over the internet or over a private network connection. Additionally, it can be scripted and automated, making it a great choice for this scenario.

AWS Documentation -- DataSync https://docs.aws.amazon.com/datasync/latest/userguide/what-is-datasync.html AWS Certified Solutions Architect Professional Official Text Book -- Chapter 4. Data Storage https://aws.amazon.com/training/learning-paths/professional-solutions-architect/

Question 3

Question Type: **MultipleChoice**

A solutions architect has an operational workload deployed on Amazon EC2 instances in an Auto Scaling group. The VPC architecture spans two Availability Zones (AZ) with a subnet in each that the Auto Scaling group is targeting. The VPC is connected to an on-premises environment and connectivity cannot be interrupted. The maximum size ol the Auto Scaling group is 20 instances in service. The VPC IPv4 addressing is as follows:

VPC CIDR: 10.0.0.0/23

AZ1 subnet CIDR: 10.0.0.0/24

AZ2 subnet CIDR: 10.0.1.0/24

Since deployment, a third AZ has become available in the Region. The solutions architect wants to adopt the new AZ without adding additional IPv4 address space and without service downtime.

Which solution will meet these requirements?

Options:

A) Update the Auto Scaling group to use the AZ2 subnet only. Delete and re-create the AZ1 subnet using hall the previous address space. Adjust the Auto Seating group to also use the new AZ1 subnet. When the instances are healthy, adjust the Auto Scaling group to use the AZ1 subnet only. Remove the current AZ2 subnet. Create a new AZ2 subnet using the second half of the address space from the original AZ1 subnet. Create a new AZ3 subnet using half the original AZ2 subnet address space, then update the Auto Scaling group to target all three new subnets.

B) Terminate the EC2 instances in the AZ1 subnet. Delete and re-create the AZ1 subnet using half the address space. Update the Auto Scaling group to use this new subnet. Repeat this for the second AZ. Define a new subnet in AZ3, then update the Auto Scaling group to target all three new subnets.

C) Create a new VPC with the same IPv4 address space and define three subnets, with one for each AZ. Update the existing Auto Scaling group to target the new subnets in the new VPC.

D) Update the Auto Scaling group to use the AZ2 subnet only. Update the AZ1 subnet to have half the previous address space. Adjust the Auto Scaling group to also use the AZ1 subnet again. When the instances are healthy, adjust the Auto Scaling group to use the AZ1 subnet only. Update the current AZ2 subnet and assign the second half of the address space from the original AZ1 subnet. Create a new AZ3 subnet using halt the original AZ2 subnet address space, then update the Auto Scaling group to target all three new subnets.

Answer:

A

Explanation:

https://aws.amazon.com/premiumsupport/knowledge-center/vpc-ip-address-range/?nc1=h_ls

It's not possible to modify the IP address range of an existing virtual private cloud (VPC) or subnet. You must delete the VPC or subnet, and then create a new VPC or subnet with your preferred CIDR block.

Question 4

Question Type: MultipleChoice

A company is migrating its three-tier web application from on-premises to the AWS Cloud. The company has the following requirements for the migration process:

* Ingest machine images from the on-premises environment.

* Synchronize changes from the on-premises environment to the AWS environment until the production cutover.

* Minimize downtime when executing the production cutover.

* Migrate the virtual machines' root volumes and data volumes.

Which solution will satisfy these requirements with minimal operational overhead?

Options:

A) Use AWS Server Migration Service (SMS) to create and launch a replication job for each tier of the application. Launch instances from the AMIs created by AWS SMS. After initial testing, perform a final replication and create new instances from the updated AMIs.

B) Create an AWS CLIVM Import/Export script to migrate each virtual machine. Schedule the script to run incrementally to maintain changes in the application. Launch instances from the AMIs created by VM Import/Export. Once testing is done, rerun the script to do a final import and launch the instances from the AMIs.

C) Use AWS Server Migration Service (SMS) to upload the operating system volumes. Use the AWS CLI import-snaps hot command 'or the data volumes. Launch instances from the AMIs created by AWS SMS and attach the data volumes to the instances. After initial testing, perform a final replication, launch new instances from the replicated AMIs. and attach the data volumes to the instances.

D) Use AWS Application Discovery Service and AWS Migration Hub to group the virtual machines as an application. Use the AWS CLI VM Import/Export script to import the virtual machines as AMIs. Schedule the script to run incrementally to maintain changes in the application. Launch instances from the AMIs. After initial testing, perform a final virtual machine import and launch new instances from the AMIs.

Answer:

A

Explanation:

SMS can handle migrating the data volumes: https://aws.amazon.com/about-aws/whats-new/2018/09/aws-server-migration-service-adds-support-for-migrating-larger-data-volumes/

Question 5

Question Type: MultipleChoice

A company's AWS architecture currently uses access keys and secret access keys stored on each instance to access AWS services. Database credentials are hard-coded on each instance. SSH keys for command-tine remote access are stored in a secured Amazon S3 bucket. The company has asked its solutions architect to improve the security posture of the architecture without adding operational complexity.

Which combination of steps should the solutions architect take to accomplish this? (Select THREE.)

Options:

A) Use Amazon EC2 instance profiles with an IAM role.

B) Use AWS Secrets Manager to store access keys and secret access keys.

C) Use AWS Systems Manager Parameter Store to store database credentials.

D) Use a secure fleet of Amazon EC2 bastion hosts (or remote access.

E) Use AWS KMS to store database credentials.

F) Use AWS Systems Manager Session Manager tor remote access

Answer:

A, C, F

Explanation:

https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager.html

Question 6

Question Type: MultipleChoice

A company is using AWS Organizations lo manage multiple accounts. Due to regulatory requirements, the company wants to restrict specific member accounts to certain AWS Regions, where they are permitted to deploy resources. The resources in the accounts must be tagged, enforced based on a group standard, and centrally managed with minimal configuration.

What should a solutions architect do to meet these requirements?

Options:

A) Create an AWS Config rule in the specific member accounts to limit Regions and apply a tag policy.

B) From the AWS Billing and Cost Management console, in the master account, disable Regions for the specific member accounts and apply a tag policy on the root.

C) Associate the specific member accounts with the root. Apply a tag policy and an SCP using conditions to limit Regions.

D) Associate the specific member accounts with a new OU. Apply a tag policy and an SCP using conditions to limit Regions.

Answer:

D

Explanation:

https://aws.amazon.com/es/blogs/mt/implement-aws-resource-tagging-strategy-using-aws-tag-policies-and-service-control-policies-scps/

Question 7

Question Type: MultipleChoice

A company wants to migrate its corporate data center from on premises to the AWS Cloud. The data center includes physical servers and VMs that use VMware and Hyper-V. An administrator needs to select the correct services to collect data (or the initial migration discovery process. The data format should be supported by AWS Migration Hub. The company also needs the ability to generate reports from the data.

Which solution meets these requirements?

Options:

A) Use the AWS Agentless Discovery Connector for data collection on physical servers and all VMs. Store the collected data in Amazon S3. Query the data with S3 Select. Generate reports by using Kibana hosted on Amazon EC2.

B) Use the AWS Application Discovery Service agent for data collection on physical servers and all VMs. Store the collected data in Amazon Elastic File System (Amazon EFS). Query the data and generate reports with Amazon Athena.

C) Use the AWS Application Discovery Service agent for data collection on physical servers and Hyper-V. Use the AWS Agentless Discovery Connector for data collection on VMware. Store the collected data in Amazon S3. Query the data with Amazon Athen a. Generate reports by using Amazon QuickSight.

D) Use the AWS Systems Manager agent for data collection on physical servers. Use the AWS Agentless Discovery Connector for data collection on all VMs. Store, query, and generate reports from the collected data by using Amazon Redshift.

Answer:

C

Explanation:

https://docs.aws.amazon.com/application-discovery/latest/userguide/discovery-agent.html

https://docs.aws.amazon.com/application-discovery/latest/userguide/discovery-connector.html

Question 8

Question Type: MultipleChoice

A large company with hundreds of AWS accounts has a newly established centralized internal process for purchasing new or modifying existing Reserved Instances. This process requires all business units that want to purchase or modify Reserved Instances to submit requests to a dedicated team for procurement or execution. Previously, business units would directly purchase or modify Reserved Instances in their own respective AWS accounts autonomously.

Which combination of steps should be taken to proactively enforce the new process in the MOST secure way possible? (Select TWO.)

Options:

A) Ensure all AWS accounts are part of an AWS Organizations structure operating in all features mode.

B) Use AWS Contig lo report on the attachment of an IAM policy that denies access to the ec2:PurchaseReservedInstancesOffering and ec2:ModifyReservedInstances actions.

C) In each AWS account, create an IAM policy with a DENY rule to the ec2:PurchaseReservedInstancesOffering and ec2:ModifyReservedInstances actions.

D) Create an SCP that contains a deny rule to the ec2:PurchaseReservedInstancesOffering and ec2: Modify Reserved Instances actions. Attach the SCP to each organizational unit (OU) of the AWS Organizations structure.

E) Ensure that all AWS accounts are part of an AWS Organizations structure operating in consolidated billing features mode.

Answer:

A, D

Explanation:

https://docs.aws.amazon.com/organizations/latest/APIReference/API_EnableAllFeatures.html

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scp-strategies.html

A: By ensuring all AWS accounts are part of an organization in AWS Organizations, it allows for centralized management and control of the accounts. This can help enforce the new purchasing process by giving a dedicated team the ability to manage and enforce policies across all accounts. D: By creating an SCP (Service Control Policy) that denies access to the ec2:PurchaseReservedInstancesOffering and ec2:ModifyReservedInstances actions, it enforces the new centralized purchasing process. Attaching the SCP to each OU (organizational unit) within the organization ensures that all business units are adhering to the new process.

Question 9

Question Type: MultipleChoice

A company standardized its method of deploying applications to AWS using AWS CodePipeline and AWS Cloud Formation. The applications are in Typescript and Python. The company has recently acquired another business that deploys applications to AWS using Python scripts.

Developers from the newly acquired company are hesitant to move their applications under CloudFormation because it would require than they learn a new domain-specific language and eliminate their access to language features, such as looping.

How can the acquired applications quickly be brought up to deployment standards while addressing the developers' concerns?

Options:

A) Create CloudFormation templates and re-use parts of the Python scripts as instance user dat
a. Use the AWS Cloud Development Kit (AWS CDK) to deploy the application using these templates. Incorporate the AWS CDK into CodePipeline and deploy the application to AWS using these templates.

B) Use a third-party resource provisioning engine inside AWS CodeBuild to standardize the deployment processes of the existing and acquired company. Orchestrate the CodeBuild job using CodePipeline.

C) Standardize on AWS OpsWorks. Integrate OpsWorks with CodePipeline. Have the developers create Chef recipes to deploy their applications on AWS.

D) Define the AWS resources using Typescript or Python. Use the AWS Cloud Development Kit (AWS CDK) to create CloudFormation templates from the developers' code, and use the AWS CDK to create CloudFormation stacks. Incorporate the AWS CDK as a CodeBuild job in CodePipeline.

Answer:

D

Explanation:

https://docs.aws.amazon.com/cdk/latest/guide/codepipeline_example.html

By using the AWS CDK, the developers can define the AWS resources using the familiar Typescript or Python programming languages, rather than learning a new domain-specific language like CloudFormation. The AWS CDK then generates the CloudFormation templates, allowing the company to standardize on CloudFormation for deployment while still leveraging the developers' expertise in Typescript or Python. The AWS CDK can be integrated as a CodeBuild job in CodePipeline, making it part of the standardized deployment process.

Question 10

Question Type: MultipleChoice

A company is planning on hosting its ecommerce platform on AWS using a multi-tier web application designed for a NoSQL database. The company plans to use the us-west-2 Region as its primary Region. The company want to ensure that copies of the application and data are available in a second Region, us-west-1, for disaster recovery. The company wants to keep the time to fail over as low as possible. Failing back to the primary Region should be possible without administrative interaction after the primary service is restored.

Which design should the solutions architect use?

Options:

A) Use AWS Cloud Formation StackSets lo create the stacks in both Regions with Auto Scaling groups for the web and application tiers. Asynchronously replicate static content between Regions using Amazon S3 cross-Region replication. Use an Amazon Route 53 DNS failover routing policy to direct users to the secondary site in us-west-1 in the event of an outage. Use Amazon DynamoDB global tables for the database tier.

B) Use AWS Cloud Formation StackSets to create the stacks in both Regions with Auto Scaling groups for the web and application tiers. Asynchronously replicate static content between Regions using Amazon S3 cross-Region replication. Use an Amazon Route 53 DNS failover routing policy to direct users to the secondary site in us-west-1 in the event of an outage. Deploy an Amazon Aurora global database for the database tier.

C) Use AWS Service Catalog to deploy the web and application servers in both Regions. Asynchronously replicate static content between the two Regions using Amazon S3 cross-Region replication. Use Amazon Route 53 health checks to identify a primary Region failure and update the public DNS entry listing to the secondary Region in the event of an outage. Use Amazon RDS for MySQL with cross-Region replication for the database tier.

D) Use AWS CloudFormation StackSets to create the stacks in both Regions using Auto Scaling groups for the web and application tiers. Asynchronously replicate static content between Regions using Amazon S3 cross-Region replication. Use Amazon CloudFront with static files in Amazon S3, and multi-Region origins for the front-end web tier. Use Amazon DynamoD8 tables in each Region with scheduled backups to Amazon S3.

Answer:

A

Explanation:

In this design, AWS Cloud Formation StackSets is used to create the stacks in both Regions, ensuring consistency across both environments. The Auto Scaling groups for the web and application tiers provide scalability and reliability, while the asynchronous replication of static content using Amazon S3 cross-Region replication ensures data availability. The use of an Amazon Route 53 DNS failover routing policy allows for fast and automatic failover to the secondary Region in the event of an outage, without the need for administrative interaction. The use of Amazon DynamoDB global tables for the database tier ensures that data is always available, even in the event of an outage.

Question 11

Question Type: MultipleChoice

A financial services company logs personally identifiable information 10 its application logs stored in Amazon S3. Due to regulatory compliance requirements, the log files must be encrypted at rest. The security team has mandated that the company's on-premises hardware security modules (HSMs) be used to generate the CMK material.

Which steps should the solutions architect take to meet these requirements?

Options:

A) Create an AWS CloudHSM cluster. Create a new CMK in AWS KMS using AWS_CloudHSM as the source (or the key material and an origin of AWS_CLOUDHSM. Enable automatic key rotation on the CMK with a duration of 1 year. Configure a bucket policy on the togging bucket thai disallows uploads of unencrypted data and requires that the encryption source be AWS KMS.

B) Provision an AWS Direct Connect connection, ensuring there is no overlap of the RFC 1918 address space between on-premises hardware and the VPCs. Configure an AWS bucket policy on the logging bucket that requires all objects to be encrypted. Configure the logging application to query the on-premises HSMs from the AWS environment for the encryption key material, and create a unique CMK for each logging event.

C) Create a CMK in AWS KMS with no key material and an origin of EXTERNAL. Import the key material generated from the on-premises HSMs into the CMK using the public key and import token provided by AWS. Configure a bucket policy on the logging bucket that disallows uploads of non-encrypted data and requires that the encryption source be AWS KMS.

D) Create a new CMK in AWS KMS with AWS-provided key material and an origin of AWS_KMS. Disable this CMK. and overwrite the key material with the key material from the on-premises HSM using the public key and import token provided by AWS. Re-enable the CMK. Enable automatic key rotation on the CMK with a duration of 1 year. Configure a bucket policy on the logging bucket that disallows uploads of non-encrypted data and requires that the encryption source be AWS KMS.

Answer:

C

Explanation:

https://aws.amazon.com/blogs/security/how-to-byok-bring-your-own-key-to-aws-kms-for-less-than-15-00-a-year-using-aws-cloudhsm/

https://docs.aws.amazon.com/kms/latest/developerguide/importing-keys-create-cmk.html

Question 12

Question Type: MultipleChoice

A solutions architect is designing a network for a new cloud deployment. Each account will need autonomy to modify route tables and make changes. Centralized and controlled egress internet connectivity is also needed. The cloud footprint is expected to grow to thousands of AWS accounts.

Which architecture will meet these requirements?

Options:

A) A centralized transit VPC with a VPN connection to a standalone VPC in each account. Outbound internet traffic will be controlled by firewall appliances.

B) A centralized shared VPC with a subnet for each account. Outbound internet traffic will controlled through a fleet of proxy servers.

C) A shared services VPC to host central assets to include a fleet of firewalls with a route to the internet. Each spoke VPC will peer to the central VPC.

D) A shared transit gateway to which each VPC will be attached. Outbound internet access will route through a fleet of VPN-attached firewalls.

Answer:

D

Explanation:

https://docs.aws.amazon.com/whitepapers/latest/building-scalable-secure-multi-vpc-network-infrastructure/centralized-egress-to-internet.html

https://docs.aws.amazon.com/whitepapers/latest/building-scalable-secure-multi-vpc-network-infrastructure/centralized-egress-to-internet.html

AWS Transit Gateway helps you design and implement networks at scale by acting as a cloud router. As your network grows, the complexity of managing incremental connections can slow you down. AWS Transit Gateway connects VPCs and on-premises networks through a central hub. This simplifies your network and puts an end to complex peering relationships -- each new connection is only made once.