- **Vendor:** Amazon
- **Exam Code:** AWS-Certified-Solutions-Architect-Professional
- **Exam Name:** AWS Certified Solutions Architect - Professional (SAP-C01)
- **New Updated Questions from Braindump2go (Updated in August/2020)**

QUESTION 643
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?

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 CLI VM 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-snapshot command for 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: B

QUESTION 644
An enterprise company's data science team wants to provide a safe, cost-effective way to provide easy access to Amazon SageMaker. The data scientists have limited AWS knowledge and need to be able to launch a Jupyter notebook instance. The notebook instance needs to have a preconfigured AWS KMS key to encrypt data at rest on the machine learning storage volume without exposing the complex setup requirements.

Which approach will allow the company to set up a self-service mechanism for the data scientists to launch Jupyter notebooks in its AWS accounts with the LEAST amount of operational overhead?

A. Create a serverless front end using a static Amazon S3 website to allow the data scientists to request a Jupyter notebook instance by filling out a form.
Use Amazon API Gateway to receive requests from the S3 website and trigger a central AWS Lambda function to make an API call to Amazon SageMaker that will launch a notebook instance with a preconfigured KMS key for the data scientists.
Then call back to the front-end website to display the URL to the notebook instance.
B. Create an AWS CloudFormation template to launch a Jupyter notebook instance using the AWS::SageMaker::NotebookInstance resource type with a preconfigured KMS key.
Add a user-friendly name to the CloudFormation template. Display the URL to the notebook using the Outputs section.
Distribute the CloudFormation template to the data scientists using a shared Amazon S3 bucket.
C. Create an AWS CloudFormation template to launch a Jupyter notebook instance using the AWS::SageMaker::NotebookInstance resource type with a preconfigured KMS key.
Simplify the parameter names, such as the instance size, by mapping them to Small, Large, and X-Large using the Mappings section in CloudFormation.
Display the URL to the notebook using the Outputs section, then upload the template into an AWS Service Catalog product in the data scientist's portfolio, and share it with the data scientist IAM role.
D. Create an AWS CLI script that the data scientists can run locally.
Provide step-by-step instructions about the parameters to be provided while executing the AWS CLI script to launch a Jupyter notebook with a preconfigured KMS key.
Distribute the CLI script to the data scientists using a shared Amazon S3 bucket.

Answer: B

QUESTION 645
A company is migrating its applications to AWS. The applications will be deployed to AWS accounts owned by business units. The company has several teams of developers who are responsible for the development and maintenance of all applications. The company is expecting rapid growth in the number of users.

The company's chief technology officer has the following requirements:

```
- Developers must launch the AWS infrastructure using AWS CloudFormation.
- Developers must not be able to create resources outside of CloudFormation.
- The solution must be able to scale to hundreds of AWS accounts.
```

Which of the following would meet these requirements? (Choose two.)

A. Using CloudFormation, create an IAM role that can be assumed by CloudFormation that has permissions to create all the resources the company needs.
Use CloudFormation StackSets to deploy this template to each AWS account.
B. In a central account, create an IAM role that can be assumed by developers, and attach a policy that allows interaction with CloudFormation.
Modify the AssumeRolePolicyDocument action to allow the IAM role to be passed to CloudFormation.
C. Using CloudFormation, create an IAM role that can be assumed by developers, and attach policies that allow interaction with and passing a role to CloudFormation.
Attach an inline policy to deny access to all other AWS services.
Use CloudFormation StackSets to deploy this template to each AWS account.
D. Using CloudFormation, create an IAM role for each developer, and attach policies that allow interaction with CloudFormation.
Use CloudFormation StackSets to deploy this template to each AWS account.

E. In a central AWS account, create an IAM role that can be assumed by CloudFormation that has permissions to create the resources the company requires.
Create a CloudFormation stack policy that allows the IAM role to manage resources.
Use CloudFormation StackSets to deploy the CloudFormation stack policy to each AWS account.

Answer: AB

QUESTION 646
A media company has a static web application that is generated programmatically. The company has a build pipeline that generates HTML content that is uploaded to an Amazon S3 bucket served by Amazon CloudFront. The build pipeline runs inside a Build Account. The S3 bucket and CloudFront distribution are in a Distribution Account. The build pipeline uploads the files to Amazon S3 using an IAM role in the Build Account. The S3 bucket has a bucket policy that only allows CloudFront to read objects using an origin access identity (OAI). During testing all attempts to access the application using the CloudFront URL result in an HTTP 403 Access Denied response.

What should a solutions architect suggest to the company to allow access the objects in Amazon S3 through CloudFront?

A. Modify the S3 upload process in the Build Account to add the bucket-owner-full-control ACL to the objects at upload.
B. Create a new cross-account IAM role in the Distribution Account with write access to the S3 bucket.
Modify the build pipeline to assume this role to upload the files to the Distribution Account.
C. Modify the S3 upload process in the Build Account to set the object owner to the Distribution Account.
D. Create a new IAM role in the Distribution Account with read access to the S3 bucket.
Configure CloudFront to use this new role as its OAI. Modify the build pipeline to assume this role when uploading files from the Build Account.

Answer: D

QUESTION 647
A company has built a high performance computing (HPC) cluster in AWS for a tightly coupled workload that generates a large number of shared files stored in Amazon EFS. The cluster was performing well when the number of Amazon EC2 instances in the cluster was 100. However, when the company increased the cluster size to 1,000 EC2 instances, overall performance was well below expectations.

Which collection of design choices should a solutions architect make to achieve the maximum performance from the HPC cluster? (Choose three.)

A. Ensure the HPC cluster is launched within a single Availability Zone.
B. Launch the EC2 instances and attach elastic network interfaces in multiples of four.
C. Select EC2 instance types with an Elastic Fabric Adapter (EFA) enabled.
D. Ensure the clusters is launched across multiple Availability Zones
E. Replace Amazon EFS win multiple Amazon EBS volumes in a RAID array.
F. Replace Amazon EFS with Amazon FSx for Lustre.

Answer: DEF

QUESTION 648
A company with multiple accounts is currently using a configuration that does not meet the following security governance policies:

```
- Prevent ingress from port 22 to any Amazon EC2 instance.
```

- Require billing and application tags for resources.
- Encrypt all Amazon EBS volumes.

A solutions architect wants to provide preventive and detective control, including notifications about a specific resource, if there are policy deviations.

Which solution should the solutions architect implement?

A. Create an AWS CodeCommit repository containing policy-compliant AWS CloudFormation templates.
Create an AWS Service Catalog portfolio.
Import the CloudFormation templates by attaching the CodeCommit repository to the portfolio.
Restrict users across all accounts to items from the AWS Service Catalog portfolio.
Use AWS Config managed rules to detect deviations from the policies.
Configure an Amazon CloudWatch Events rule for deviations, and associate a CloudWatch alarm to send notifications when the TriggeredRules metric is greater than zero.

B. Use AWS Service Catalog to build a portfolio with products that are in compliance with the governance polices in a central account.
Restrict users across all accounts to AWS Service Catalog products.
Share a compliant portfolio to other accounts. Use AWS Config managed rules to detect deviations from the policies.
Configure an Amazon CloudWatch Events rule to send a notification when a deviation occurs.

C. Implement policy-compliant AWS CloudFormation templates for each account, and ensure that all provisioning is completed by CloudFormation.
Configure Amazon Inspector to perform regular checks against resources.
Perform policy validation and write the assessment output to Amazon CloudWatch Logs.
Create a CloudWatch Logs metric filter to increment a metric when a deviation occurs.
Configure a CloudWatch alarm to send notifications when the configured metric is greater than zero.

D. Restrict users and enforce least privilege access using AWS IAM.
Consolidate all AWS CloudTrail logs into a single account.
Send the CloudTrail logs to Amazon Elasticsearch Service (Amazon ES).
Implement monitoring, alerting, and reporting using the Kibana dashboard in Amazon ES and with Amazon SNS.

Answer: C

QUESTION 649
A company is manually deploying its application to production and wants to move to a more mature deployment pattern. The company has asked a solutions architect to design a solution that leverages its current Chef tools and knowledge. The application must be deployed to a staging environment for testing and verification before being deployed to production. Any new deployment must be rolled back in 5 minutes if errors are discovered after a deployment.

Which AWS service and deployment pattern should the solutions architect use to meet these requirements?

A. Use AWS Elastic Beanstalk and deploy the application using a rolling update deployment strategy.
B. Use AWS CodePipeline and deploy the application using a rolling update deployment strategy.
C. Use AWS CodeBuild and deploy the application using a canary deployment strategy.
D. Use AWS OpsWorks and deploy the application using a blue/green deployment strategy.

Answer: A

QUESTION 650

A company has been using a third-party provider for its content delivery network and recently decided to switch to Amazon CloudFront. The development team wants to maximize performance for the global user base. The company uses a content management system (CMS) that serves both static and dynamic content. The CMS is behind an Application Load Balancer (ALB) which is set as the default origin for the distribution. Static assets are served from an Amazon S3 bucket. The Origin Access Identity (OAI) was created property and the S3 bucket policy has been updated to allow the GetObject action from the OAI, but static assets are receiving a 404 error.

Which combination of stops should the solutions architect take to fix the error? (Choose two.)

A. Add another origin to the CloudFront distribution for the state asses.
B. Add a path-based rule to the ALB to forward requests for the static assets.
C. Add an RTMP distribution to allow caching of both static and dynamic content.
D. Add a behavior to the CloudFront distribution for the path pattern and the origin of the static assets.
E. Add a host header condition to the ALB listener and forward the header from CloudFront to add traffic to the allow list.

Answer: AB

QUESTION 651
A financial services company logs personally identifiable information to 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?

A. Create an AWS CloudHSM cluster. Create a new CMK in AWS KMS using AWS_CloudHSM as the source for 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 logging bucket that 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: D

QUESTION 652
A solutions architect is implementing infrastructure as code for a two-tier web application in an AWS CloudFormation template. The web frontend application will be deployed on Amazon EC2 instances in an Auto Scaling group. The backend database will be an Amazon RDS for MySQL DB instance. The database password will be rotated every 60 days.

How can the solutions architect MOST securely manage the configuration of the application's database credentials?

A. Provide the database password as a parameter in the CloudFormation template.
Create an initialization script in the Auto Scaling group's launch configuration UserData property to reference the password parameter using the Ref intrinsic function.
Store the password on the EC2 instances.
Reference the parameter for the value of the MasterUserPassword property in the AWS::RDS::DBInstance resource using the Ref intrinsic function.

B. Create a new AWS Secrets Manager secret resource in the CloudFormation template to be used as the database password.
Configure the application to retrieve the password from Secrets Manager when needed.
Reference the secret resource for the value of the MasterUserPassword property in the AWS::RDS::DBInstance resource using a dynamic reference.

C. Create a new AWS Secrets Manager secret resource in the CloudFormation template to be used as the database password.
Create an initialization script in the Auto Scaling group's launch configuration UserData property to reference the secret resource using the Ref intrinsic function.
Reference the secret resource for the value of the MasterUserPassword property in the AWS::RDS::DBInstance resource using the Ref intrinsic function.

D. Create a new AWS Systems Manager Parameter Store parameter in the CloudFormation template to be used as the database password.
Create an initialization script in the Auto Scaling group's launch configuration UserData property to reference the parameter.
Reference the parameter for the value of the MasterUserPassword property in the AWS::RDS::DBInstance resource using the Fn::GetAtt intrinsic function.

Answer: D

QUESTION 653
A company built an application based on AWS Lambda deployed in an AWS CloudFormation stack. The last production release of the web application introduced an issue that resulted in an outage lasting several minutes. A solutions architect must adjust the deployment process to support a canary release.

Which solution will meet these requirements?

A. Create an alias for every new deployed version of the Lambda function.
Use the AWS CLI update-alias command with the routing-config parameter to distribute the load.

B. Deploy the application into a new CloudFormation stack.
Use an Amazon Route 53 weighted routing policy to distribute the load.

C. Create a version for every new deployed Lambda function.
Use the AWS CLI update-function- configuration command with the routing-config parameter to distribute the load.

D. Configure AWS CodeDeploy and use CodeDeployDefault.OneAtATime in the Deployment configuration to distribute the load.

Answer: C