Practice Questions

1. There is an application, which consists of the EC2 instance in an Auto-scaling group. During a particular timeframe, users are complaining about the poor response time of application due to the increase in traffic. You have to deploy a new EC2 instance for an Auto-scaling group when utilization of CPU is greater than 60% for two consecutive periods of 5 minutes. How will you do this?

    A. By decreasing the consecutive number of collection periods
    B. By increasing the minimum number of instances in the Auto-scaling group
    C. By decreasing the collection period to ten minutes
    D. By decreasing the threshold CPU utilization percentage at which to deploy a new instance

Answer: B

Explanation: If you increase the minimum number of instances, the application will be running even when the load on the website is not high.

2. You have a setup in AWS that contains an elastic load balancer, an Auto-scaling group, which launches EC2 instances, and AMIs with your code pre-installed. You want a cost-effective solution and deploy the updates up to a certain number of users from which you are able to revert quickly. What is the best approach?

    A. Create a second ELB and a new Auto-scaling group assigned a new launch configuration. Create a new AMI with an updated app. Use Route53 weighted round robin records to adjust the proportion of traffic hitting the two ELBs

    B. Create a new AMI's with a new app. Then use the new EC2 instances in half proportion to the older instances
    C. Redeploy with AWS Elastic Beanstalk and Elastic Beanstalk versions. Use Route53 weighted round robin records to adjust the proportion of traffic hitting the two ELBs
    D. Create a full second stack of instances, cut the DNS over to the new stack of instances, and change the DNS back if a rollback

is needed

Answer: A

Explanation: The weighted routing policy of Route53 is used to direct the proportion of traffic to your application. The best policy is to create new ELB, attach to the Auto-scaling group and then divert the traffic by using Route53.

3. A specific process running an application, which is critical to application functionality and health check is added to the Auto-scaling group. The instances are showing healthy, but the application is not working properly. What is the problem with the health checks?

    A. You do not have the time range in the health check properly configured

    B. It is not possible for a health check to monitor a process that involves an application
    C. The health check is not configured properly
    D. The health check is not checking the application process

Answer: D

Explanation: In the case of the custom health check, you can send information from your health check to Auto-scaling. So that Auto-scaling can utilize this information. For example, if you determine that an instance is not functioning as expected, you can set the health status of the instance to “Unhealthy”. The next time that Auto-scaling performs a health check on the instance, it will determine that the instance is unhealthy and then launch a replacement instance.

4. In the AWS ap-south-1 region, Alan developed an AWS console git repository called "MyRepo". The required credentials for cloning of the repository were configured in CodeCommit. Alan is now developing the software with a Ubuntu Linux machine. Which commands can the repository be cloned in a local directory from CodeCommit? (Choose 2)

    A. git clone ssh://git-codecommit.ap-south-1.amazonaws.com/v1/repos/MyDemoRepo my-demo-repo

B. git clone ssh://git-codecommit.aws.com/v1/repos/MyDemoRepo my-demo-repo
C. git clone http://git-codecommit.ap-south-1.amazonaws.com/v1/repos/MyDemoRepo my-demo-repo
D. git clone https://git-codecommit.ap-south-1.amazonaws.com/v1/repos/MyDemoRepo my-demo-repo
E. git clone https://git-codecommit.amazonaws.com/v1/repos/MyDemoRepo my-demo-repo

Answer: A and D

Explanation: You must have a clone URL for the relevant CodeCommit repository to clone a CodeCommit repository. One thing to note is that the repository name and AWS Region must be included in this URL. The best way to get the URL is to click on 'clone URL'.

By using SSH and HTTPS, you can connect to the repository.

5. A new employee accidentally removed a CodeCommit repository branch last week. For certain key repositories, when someone deletes any branch in CodeCommit, Julie is told to create a notification. In CodeCommit, you plan to set a trigger. Which service can be used for this purpose?

    A. For the repository event of deletion of branches, create a trigger in CodeCommit to an AWS CloudWatch Event to provide warnings

    B. For the repository event of deletion of branches, create a trigger in CodeCommit to an Amazon SNS topic to provide notifications to users
    C. In CodeCommit, create a trigger for branch deletion event to an Amazon SQS queue to provide warnings to users
    D. In CodeCommit, create a trigger to a Lambda function, which works with AWS Simple Email Service to send emails notifying users when a branch is deleted

Answer: B

Explanation: You can set triggers for a repository of CodeCommit to move

code or to trigger actions for other events. SNS and Lambda are the supported trigger services. For each CodeCommit repository, you can create up to 10 triggers, as SNS is the easiest method. Users simply have to subscribe to SNS topic and be notified.

6. Multiple applications are running on AWS. The company wants you to develop a tool that immediately calls the team when the alarm is triggered. You also have to make sure that alarm for the on-call team generated must handle to notify the correct team at the correct time. What steps should be taken to implement this?

    A. Create an Amazon SNS topic and an Amazon SQS queue. Configure the Amazon SQS queue as a subscriber to the Amazon SNS topic. Configure CloudWatch alarms to notify this topic when an alarm is triggered. Create an Amazon EC2 Auto-scaling group with both minimum and desired instances configured to 0. The worker node in this group spawns when messages are added to the queue. Workers then use Amazon Simple Email Service to send messages to your on-call teams

    B. Create an Amazon SNS topic and configure your on-call team email addresses as a subscriber. Use the AWS SDK tools to integrate your application with Amazon SNS and send messages to this new topic. Notifications will be sent to on-call users when a CloudWatch alarm is triggered

    C. Create an Amazon SNS topic and configure your on-call email addresses as subscribers. Create a secondary Amazon SNS topic for alarms and configure your CloudWatch alarms to notify this topic when triggered. Create an HTTP subscriber to this topic that notifies your application via HTTP post when an alarm is triggered. Use the AWS SDK tools to integrate your application with Amazon SNS and messages to the first topic so that on-call engineers receive alerts

    D. Create an Amazon SNS topic for each on-call group, and configure each of these with the team member emails as a subscriber. Create another Amazon SNS topic and configure your CloudWatch alarms to notify this topic when triggered. Create an HTTP subscriber to this topic that notifies your

application via HTTP post when an alarm is triggered. Use the AWS SDK tools to integrate your application with Amazon SNS and send messages to the correct team topic when on shift

Answer: D

Explanation: This option fulfills all the requirements. First step is to create an SNS group so that specific member gets the email address and ensure that the application uses HTTP endpoint and SDK for publishing messages.

You cannot use SQS service for this purpose, and as the message notification needs to be given to only specific members and not to all members so other options are also not valid.

7. While analyzing the metrics, you got to know that the company website is experiencing response times higher than anticipated during peak hours. You rely on Auto-scaling to make sure you are scaling your time during peak windows. What can you do to enhance your Auto-scaling policy to decrease the response time, which is high?

    A. Push custom metrics to CloudWatch to monitor your CPU and network bandwidth from your servers, which will allow your Auto-scaling policy you have better fine-grain insight

    B. Increase your Auto-scaling group's number of max servers
    C. Create a script that runs and monitors your servers; when it detects an anomaly in load, it posts to an Amazon SNS topic that triggers elastic load balancing to add more servers to the load balancer
    D. Push custom metrics to CloudWatch for your application that include more detailed information about your web application, such as how many requests it is handling and how many are waiting to be processed

Answer: B and D

Explanation: Option B is valid because the max server is low. Therefore, the application cannot handle the peak load.

Option D ensures that Auto-scaling can scale the group to the right metrics.

8. An organization has an application that uses stateless web tier on

EC2 instances that are behind the ELB and use RDS read replicas. From the following options, which option is best to implement self-healing and cost effective architecture?

A. Use a larger Amazon EC2 instance type for the web server tier and a larger DB instance type for the data storage layer to ensure that they do not become unhealthy

B. Use an Amazon RDS Multi-AZ deployment
C. Set up a third-party monitoring solution on a cluster of Amazon EC2 instances in order to emit custom CloudWatch metrics to trigger the termination of unhealthy Amazon EC2 instances
D. Set up an Auto-scaling group for the web server tier along with an Auto-scaling policy that uses the Amazon EC2 CPU utilization CloudWatch metric to scale the instances

Answer: B and D

Explanation: Amazon RDS Multi-AZ deployments provide enhanced DB (database) availability and durability to make them natural for production database workloads. Amazon RDS creates automatically a primary DB Instance when providing a multi-AZ DB Instance and synchronizes the data to a standby instance in a different Availability Zone (AZ). For scaling of EC2 instance in Auto-scaling, you must use the CPU utilization metric of the instance.

9. There is a code repository that is stored in Amazon S3. In a recent audit of security control, some questions arose about maintaining the integrity of the data stored in Amazon S3 and securely deploying codes from S3 to running applications on EC2 in a virtual private cloud. What can you do to reduce these concerns? (Choose 2)

A. Add an Amazon S3 bucket policy with a condition statement to allow access only from Amazon EC2 instances with RFC 1918 IP addresses and enable bucket versioning

B. Use a configuration management service to deploy AWS identity and access management user credentials to the Amazon EC2 instances. Use these credentials to securely

access the Amazon S3 bucket when deploying code

C. Create an Amazon identity access and management role with authorization to access the Amazon S3 bucket, and launch all of your application's Amazon EC2 instances with this role
D. Use AWS data pipeline to lifecycle the data in your Amazon S3 bucket to Amazon Glacier on a weekly basis
E. Use AWS pipe line with multi-factor authentication to securely deploy code from the Amazon S3 bucket to your Amazon EC2 instances
F. Add an Amazon S3 bucket policy with a condition statement that requires multi-factor authentication in order to delete objects and enable bucket versioning

Answer: C and F

Explanation: MFA delete is enabled on versioning bucket by adding another layer of protection. In order to perform the permanent deletion, you need to provide the AWS account's access key and code from the MFA device.

IAM functions are built to safely request APIs from your applications, without you having to manage the security credentials of the applications. You can delegate authorization to make API requests using IAM roles instead of creating and distributing your AWS credentials.

10. The operation and development department wants a place where it can show both the operation system and application logs. What should you do to activate this service using AWS? (Choose 2)

A. Using AWS CloudFormation, create a CloudWatch Log. Log groups and send the operating system and application logs of interest using the CloudWatch logs agent

B. Using AWS CloudWatch and configuration management, set up remote logging to send events through UDP packets to CloudTrail
C. Using configuration management, set up remote logging to send events to Amazon Kinesis and insert these into Amazon cloud search or Amazon RedShift, depending on the available analytic tool
D. Using AWS CloudFormation, merge the application logs with

the operating system logs, and use IAM roles to allow both teams to have access to view console output from Amazon EC2

Answer: A and C

Explanation: Amazon CloudWatch logs can be used to monitor, store, and access log files from Amazon Elastic Compute Cloud (Amazon EC2), AWS Cloud Trail instances, and other sources. The related log data can then be collected from CloudWatch Logs.

11. An enterprise is using CloudFormation for its application deployment and now they want a cost-effective solution for creating a rolling deployment with minimum downtime. How can this be executed? (Choose 2)

A. By re-deploying the application using a CloudFormation template to deploy Elastic Beanstalk

B. By re-deploying with a CloudFormation template and defining update policies on Auto-scaling groups in the template
C. By using the AutoScalingRollingUpdate attribute to specify how CloudFormation handles updates to Auto-scaling Group Resource
D. By tearing down the old stack after each stack is deployed

Answer: B and C

Explanation: The Auto-scaling group resource supports an UpdatePolicy attribute. This defines how an Auto-scaling group is updated when an update in CloudFormation occurs. The Auto-scaling group is updated by the rolling update, which is performed by specifying the AutoscalingRollingUpdate Policy. This retains the same Auto-scaling group by replacing the old instances with a new one.

12. John created a mobile application whose main purpose is photo sharing. In the very beginning, there will be approx. ten thousand users expected. He uses AWS S3 to store images and has to determine how users can be authenticated in order to access the images. The storage of these images also needs to be managed. Which steps should John take next? (Choose 2)

A. Use a key-based naming scheme comprised from the user ID's for all user objects in a single Amazon S3 bucket

B. Create an Amazon S3 bucket per user, and use the application to generate the S3 URL for the appropriate content
C. Use AWS IAM user accounts as the application level user database, and offload the burden of authentication from the application code
D. Authenticate users at the application level, and use AWS security token service (STS) to grant token-based authorization to S3-object
E. Authenticate users at the application level, and send an SMS token message to the user. Create an Amazon S3 bucket with the same name as the SMS message token, and move the user's objects to that bucket

Answer: A and D

Explanation: The AWS STS is the web service that enables you to request temporary and limited access credentials to IAM users and for users that you authenticate. This token is then used for accessing an object in S3.

13. Michael uses an Auto-scaling group for its application. The Auto-scaling group has 2 AZs, one AZ contains 4 EC2 instances and another AZ contains 3 instances. None of the instances are protected from the scale. What is going to happen, depending on the Auto-scaling termination policy set by default?

A. Auto-scaling will select an instance to terminate randomly
B. Auto-scaling will select the AZ with 4 EC2 instances and terminate an instance

C. Auto-scaling will terminate un-protected instances in the availability zone with the oldest launch configuration
D. Auto-scaling will terminate all unprotected instances that are the closest to the next billing hour

Answer: B

Explanation: To ensure that the network architecture spans availability

zones evenly, the default termination policy is designed. When using the default termination policy, Auto-scaling selects an instance to terminate as follow:

Auto-scaling determines if instances exist in multiple AZs. If so, you must select the AZ with the most instances and at least one instance that is not protected from the scale. In case of more than one availability zone having the same number of instances, Auto-scaling selects the availability zone with the instances with the oldest configuration.

14. David manages a continual integration application to monitor version control for changes and then run a complete suite of build tests on new Amazon EC2 instances. What should he do to guarantee the lowest overall cost while running as many tests as possible in parallel?

A. Perform syntax checking on the continuous integration system before launching a new Amazon EC2 instance for build test, unit, and integration tests

B. Perform all test on continuous integration system, using AWS OpsWorks for units, integration, and build tests
C. Perform syntax and build test on the continuous integration system before launching the new Amazon EC2 instance units and integration tests
D. Perform syntax checking on the continuous integration system before launching a new AWS data pipeline for coordinating the output of unit, integration, and build tests

Answer: C

Explanation: When developers want to integrate the code to the shared repository for as much time as they want, then Continuous Integration is the best practice tool for development. Each of the check-ins is verified by an automated build, which allows the team to detect problems early.

15. There is an Auto-scaling group associated with a load balancer. You want to suspend Auto-scaling AddToLoadBalancer for some time. What is the effect of this on launched instances during the suspension period?

A. The instances will be registered with ELB once the process has resumed
B. The instances will not be registered with ELB. You must manually register when the process is resumed

C. Auto-scaling will not launch the instances during this period because of the suspension
D. It is not possible to suspend the AddToLoadBalancer process

Answer: C

Explanation: Auto-scaling starts up the instances but does not add them to load balancers or target groups until you suspend the AddToLoadBalancer. Auto-scaling will resume adding instances to load balancer or target group if you resume AddToLoadBalancer. Auto-scaling, however, does not include the instances, which are launched during that process. You have to manually add them.

16. You decided to change the instance type of your production instances that run in the Auto-scaling group. To launch our architecture, we used the CloudFormation template and currently used 4 instances in production. The service cannot be interrupted; therefore, two instances should always run during the update. Which of the following options can be applicable?

A. Auto-scalingReplacingUpdate

B. Auto-scalingScheduledAction
C. Auto-scalingRollingUpdate
D. Auto-scalingIntegrationUpdate

Answer: C

Explanation: The UpdatePolicy attribute is provided by the AWS::Auto-scaling::Auto-scalingGroup asset. This is to describe the updating of the Auto-scaling resource group when an update to the CloudFormation stack takes place. A common approach is to perform a rolling update for updating an Auto-scaling group by defining the Auto-scalingRollingUpdate policy. This keeps the same Auto-scaling Group and, according to the specified parameters, replaces old instances with new ones.

17. A company wants to use Blue/Green Deployment for its application, which runs behind a load balancer on Amazon EC2 instances. How this be done for every deployment?

A. By creating a new load balancer with new Amazon EC2 instances, carrying out the deployment, and then switching DNS over to the new load balancer using Amazon Route53 after testing

B. By setting up Amazon Route53 health checks to fail over from any Amazon EC2 instance that is currently being deployed to
C. By creating a test stack for validating the code using AWS CloudFormation, and then deploying the code to each production Amazon EC2 instance
D. By launching more Amazon EC2 instance to ensure high availability, de- registering each Amazon EC2 instance from the load balancer, then upgrading, testing, and registering it again with the load balancer

Answer: A

Explanation: For blue/green deployment you need to create new ELB, which needs to be used for pointing to new production changes. To assign traffic to two ELBs on the basis of 80% to 20% traffic, use the Weighted Routing Policy for Route53. Change the percentage value as per requirement. Once the changes are tested, the traffic will route 100% via Route53 to the new ELB. For the figure you can visit:
https://d0.awsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf

18. An enterprise has its application, which is hosted on EC2 instances behind ELB. Now, some errors are observed in the application. In order to diagnose the error, the enterprise wants to check the ELB access logs but they are empty. What might be the reason behind this?

A. The enterprise does not have the appropriate permissions to access the logs

B. The enterprise does not have the CloudWatch metrics

correctly configured

C. Access logging is an optional feature of ELB that is disabled by default
D. ELB access logs are only available for the maximum of one week

Answer: C

Explanation: ELB provides an access log that captured the detailed information about the request sent to load balancer. Access logging is disabled by default but can be enabled, it captures the logs and stores them in the S3 bucket. That log contains information such as the requester's IP, latencies, request paths and server responses at the time the request was sent. You can use these logs to evaluate patterns of traffic and to resolve problems.

19. Harry stores sensitive information on an EBS volume attached to EC2 instance in his application. What can he do to protect this data? (Choose 2)

A. Unmount the EBS volume, take a snapshot and encrypt it. Re-mount the Amazon EBS volume
B. Copy the unencrypted snapshot and check the box to encrypt the new snapshot. Volumes restored from this encrypted snapshot will also be encrypted

C. It is not possible to encrypt an EBS volume, you use a lifecycle policy to transfer data to S3 for encryption
D. Create and mount a new encrypted Amazon EBS volume. Move the data to the new volume. Delete the old Amazon EBS volume

Answer: B and D

Explanation: By following these two methods given below, he can protect his data:

- To migrate data from encrypted to un-encrypted volume, create a destination volume and attach the destination volume to the instance that hosts the data to migrate. Copy the data from the source directory to the destination volume. For copying, use bulk-

copy utility

- Create a snapshot of the unencrypted volume and copy snapshots with an encrypted parameter. Now, restore the encrypted snapshot to the new encrypted volume

20. Company ABC works on multiple projects for different clients, the development environment of each project is different and requires multiple OS, tools, programming languages, and runtime. Now, the company is planning to use AWS CodeBuild to build an artifact for each client. From the following, which option is not a supported Docker platform for AWS CodeBuild?

A. Redhat OS

B. Amazon Linux 2
C. Windows Server Core 2016
D. Ubuntu 18.04

Answer: A

Explanation: Redhat OS is not supported in AWS CodeBuild. Amazon Linux 2, Windows Server Core 2016 and Ubuntu 18.04 are supported platforms in AWS CodeBuild.

21. There is a large multi-tiered Windows-based web application situated behind a load balancer, running on an EC2 instance. The problem of slow customer page load time occurs, and your manager asks you to sort this problem out. You must ensure that customer load time is not affected by too many requests per second. Which one of the following techniques should you use to solve this problem?

A. Re-deploy your infrastructure using the AWS CloudFormation template. Configure Elastic load balancing health check to initiate a new AWS CloudFormation stack when health checks return fail

B. Re-deploy your infrastructure using an AWS CloudFormation template. Spin up a second AWS CloudFormation stack. Configure Elastic load balancing spillover functionality to spill

over any slow connections to the second AWS CloudFormation stack

C. Re-deploy your application using an Auto-scaling template. Configure the Auto-scaling template to spin up a new Elastic Beanstalk application when the customer load time surpasses your threshold
D. Re-deploy your infrastructure using AWS CloudFormation, Elastic Beanstalk, and Auto-scaling. Set up your Auto-scaling group policies to scale based on the number of requests per second as well as the current customer load time

Answer: D

Explanation: Auto-scaling is used to ensure that there will be a valid number of EC2 instances, which can handle the load of the application. The desired number of instances can be specified at the time of the creation of the group or thereafter. In Auto-scaling group, you need to define both minimum and a maximum number of instances as per the requirement of your application. Auto-scaling can launch or terminate the instances as the demand for your application increases or decreases if you specify scaling policies.

Options A and B are invalid because Auto-scaling is required for an application to be able to handle the traffic. Option C is also invalid because there is no Auto-scaling template.

22. A company is concerned with the extreme increasing cost as its management has reported the increase in monthly bills. After reviewing the billing report, it is noticed that there is an increase in data transfer costs. Now, how can Alan provide management with better insight into data transfer use?

A. Update Amazon CloudWatch metrics to use 5-second granularity, which will give better-detailed metrics that can be combined with the billing data to pinpoint anomalies
B. Deliver custom metrics to Amazon CloudWatch per application that breaks down application data transfer into multiple, more specific data points
C. Use Amazon CloudWatch logs to run a map-reduce on logs to determine high usage and data transfer
D. Using Amazon CloudWatch metrics, pull elastic load balancing

outbound data transfer metrics monthly, and include them with their billing report to show, which application is causing higher bandwidth usage

Answer: B

Explanation: The custom metrics can be published to CloudWatch using CLI and API and can be viewed statistically with the AWS management console. If you have custom metrics specific to your application, you can give a breakdown to the management on the exact issue.

23. Currently, an infrastructure is running on EC2 instance behind an Auto-scaling group. The application logs are written on ephemeral storage. The company experienced a major bug, which triggered the Auto-scaling group up and down before successfully retrieving the logs of the server. What technique is suitable to retrieve logs?

A. Configure the ephemeral policies on your Auto-scaling group to back up on terminate

B. Configure your Auto-scaling policies to create a snapshot of all ephemeral storage on terminate
C. Install the CloudWatch logs agent on your AMI, and configure CloudWatch logs agent to stream your logs
D. Install the CloudWatch monitoring agent on your AMI, and set up a new SNS alert for CloudWatch metrics that trigger the CloudWatch monitoring agent to backup all logs on the ephemeral drive

Answer: C

Explanation: CloudWatch logs are used to monitor applications and systems using log data. It can track the number of errors that occurred in the application and send a notification when the errors are exceeded from the specified threshold value. CloudWatch log uses log data, therefore, no changes in code are required.

24. Susan has a project in which she needs to deploy an infrastructure on the AWS CloudFormation template. The infrastructure supports multi-tier applications. She needs to perform

the task of organizing AWS CloudFormation resources for the future so that different departments such as Networking and Security can review the architecture before it goes to Production. How should she use the already existing workflows to accommodate each department?

A. Separate the AWS CloudFormation template into a nested structure that has individual templates for the resources that are to be governed by different departments, and use the outputs from the networking and security stacks for the application template that you control
B. Organize the AWS CloudFormation template so that related resources are next to each other in the template, such as VPC subnets and routing rules for networking and security groups and IAM information for security
C. Organize the AWS CloudFormation templates, so that related resources are next to each other in the template for each department's use, leverage your existing continuous integration tool to constantly deploy changes from all parties to the production environment, and then run tests for validation
D. Use a custom application and the AWS SDK to replicate the resources defined in the current AWS CloudFormation template, and use the existing code review system to allow other departments to approve changes before altering the application for future deployments.

Answer: A

Explanation: When your infrastructure is growing, common patterns emerge in each template, where you declare the same components. You can separately create templates for these common components. Thus, you can mix and match the various templates and you can also create a single unified stack using nested stacks. Stacks that are nested are stacks that create additional stacks. Use AWS::CloudFormation::Stackresource to link other templates in your template to build nested stacks.

25. Peter has a team for developing some Java applications for micro-services. AWS CodeBuild is responsible for the construction of the CI / CD pipeline and supplies S3 storage artefacts. The team

manages multiple build-specific files instead of adding explicitly the build commands when the build is running. What description is INCORRECT about the use of buildspec files?

A. Build spec files must be expressed in YAML format

B. There should be only one build spec file for a given build project
C. The build spec file can be placed in any location under the top level directory
D. Users can override the default build spec file name such as buildspec_test.yml

Answer: C

Explanation: AWS CodeBuild service has used buildspec files as a collection of build commands and related settings. So buildspec files are stored in root level directory and cannot be placed at any other location.

26. An organization has messages in the SQS queue, which are processed by instances placed in an Auto-scaling group. Depending upon the queue size, the group performs scaling. In the processing, there is a third-party service calling, which tells about the failed and repeated calls from their side. The organization found that instances were terminated as the team scales during processing. What economic solution can they use to reduce the number of incomplete trials?

A. Create a new Auto-scaling group with minimum and maximum of 2s and instances running web proxy software. Configure the VPC route table to route HTTP traffic to these web proxies
B. Modify the application running on the instances to put itself into an Auto-scaling standby state while it processes a task and return itself to ‘Inservice’ when the processing is complete

C. Modify the application running on the instances to enable termination protection while it processes a task and disables it when the processing is complete
D. Increase the minimum and maximum size for the Auto-scaling

group. Change the scaling policies, so they scale less dynamically

Answer: B

Explanation: You can put the instances into a standby state. After the processing is completed, the instances come back to the state where the Auto-scaling group governs them, in order to avoid termination of processing instances.

27. To maintain the version control and achieve automation for the application in your organization, you are requested to use CloudFormation. How can you best maintain multiple environments while keeping the cost down while using CloudFormation?

A. By using CloudFormation custom resources to handle dependencies between stacks

B. By creating multiple templates in one CloudFormation stack
C. By combining all resources into one template for version control and automation
D. By creating separate templates based on functionality and nested stacks with CloudFormation

Answer: D

Explanation: When your infrastructure is growing, common patterns emerge in each template, where you declare the same components. You can separately create templates for these common components. Thus, you can mix and match the various templates and you can also create a single unified stack using nested stacks. Stacks that are nested are stacks that create additional stacks. Use AWS::CloudFormation::Stackresource to link other templates in your template to build nested stacks.

28. Martin works as a DevOps engineer for a company focused on artificial intelligence services. As part of the AI products, several functions of Lambda are being developed. He is responsible for the development of Lambda through AWS CodeDeploy. And a new build is waiting for online deployment every week. There are a number of steps for

CodeDeploy:

- Specifying an AppSpec file, which contains instructions including the Lambda functions to be deployed

- Creating a deployment, which is the process of installing contents
- Creating a CodeDeploy application
- Specifying a deployment group for settings and configurations

Which sequence does he follow for proper configuration?

A. 3-4-1-2

B. 1-2-3-4
C. 4-3-2-1
D. 1-3-2-4

Answer: A
Explanation: In order to use AWS CodeDeploy, first create an application, specify the deployment group and deployment configuration then specify the Appspec file. In the end, create a deployment.

29. Robert is using an application, which performs workflow and operation but takes a long time for completing. Which service in Elastic Beanstalk environment can be used to perform this task

A. Manages the ELB and runs a daemon process on each instance

B. Manages Lambda functions and runs a daemon process on each instance
C. Manages an Amazon SQS queue and runs a daemon process on each instance
D. Manages an Amazon SNS Topic and runs a daemon process on each instance

Answer: C
Explanation: Elastic Beanstalk makes this process easier through the Amazon SQS queue management and the execution of a daemon system for

each instance reading for you from the queue. The HTTP POST request is sent to http:/localhost/, with the content of the queue in the context, when the daemon pulls an object from the queue. The long run job in response to POST is what the application has to do.

30. John runs an online store with Elastic Beanstalk. The store is based on an e-commerce open source platform and is installed in an auto-scaling group in several instances. For the online store, your development team will generate new "extensions". These extensions include both PHP source code and an SQL upgrade script to update the database schema. He has noticed that there is an error in the extension deployment when the SQL upgrade script is running. After further investigation, he realized that the SQL script is executed in all Amazon EC2 instances. Now, he has to make sure that the SQL script is only executed once per deployment regardless of instances running at that time. What can he do to achieve this?

A. Make use of the Amazon EC2 metadata service to query whether the instance is marked as the leader in the Auto-scaling group. Only execute the script if “true” is returned

B. Use a “Container Command” within an Elastic Beanstalk configuration file to execute the script, ensuring that the “leader only” flag is set to true
C. Use a “Solo Command” within an Elastic Beanstalk configuration file to execute the script. The Elastic Beanstalk service will ensure that the command is only executed once
D. Update the Amazon RDS security group to only allow write access from a single instance in the Auto-scaling group; that way, only one instance will successfully execute the script on the database

Answer: B

Explanation: Container Command runs after the application and web server have been setup, and the application version archive has been extracted before the application version is deployed. The container-commands key is used to execute commands that affect the application source code. Before the application's source code is extracted, non-container commands and other

customization operations are carried out.
You can use 'leader_only' to execute the command on one instance or to set up a test to only execute this command if a test command is valid. Leader-only container commands are only executed during environment creation and deployments, while other commands and server customization operations are performed every time an instance is provisioned or updated. Leader-only container commands are not executed due to launch configuration changes, such as a change in the AMI ID or instance type.

31. A number of AWS products have already been used by your client. Nonetheless, Jenkins installs on local servers in most of the current pipelines. The team will move Jenkins to the AWS CodePipeline in the next quarter. All new pipelines should be of three stages: source, build and deployment. In which combinations of CodePipeline services can be used to form a new pipeline? (Choose 3)

A. Source(S3) - Build(CodeBuild) - Deploy(CodeDeploy)

B. Source(GitHub) - Build(ECS) - Deploy(Elastic Beanstalk)
C. Source(CodeCommit) - Build(CodeBuild) - Deploy(CloudFormation)
D. Source(Bitbucket) - Build(Jenkins) - Deploy(S3)
E. Source(ECR) - Build(Jenkins) - Deploy(OpsWorks)

Answer: A, C, and E
Explanation: When a new pipeline is created, you need to configure the source, build and deploy stages:

- Source stage includes- CodeCommit, ECR, GitHub, and S3
- Build stage includes- Jenkins or CodeBuild
- Deploy stage includes- CodeDeploy, ElasticBeanstalk, CloudFormation, S3, ECS and Service Catalog

32. John has a CloudFormation template in AWS. Now he wants to change the alarm threshold defined in the template under CloudWatch Alarm. How can he achieve this?

A. By deleting the current clouformation template and creating a new one that will update the current resources

B. By updating the template and then updating the stack with the new template. Only those resources that need to be changed will be changed. All other resources, which do not need to be changed will remain as they are
C. By updating the template and then updating the stack with the new template. Automatically all resources will be changed in the stack
D. Currently, there is no option to change what is already defined in Cloudformation templates

Answer: B

Explanation: In CloudFormation, whenever you perform changes like settings or resources, the stack is updated with a new template and changes perform to only those resource on which changes occurred. Otherwise, remaining resources work same as they are in the previous template.

33. In the context of your ongoing application, an I / O load performance test is performed before it is deployed to production with new AMIs. The app is running with one EBS PIOPS volume per instance and requires consistent I / O performance. To ensure that I / O load performance tests produce the correct results repeatedly, which of the following option must be carried out?

A. Ensure that the I/O block sizes for the test are randomly selected
B. Ensure that snapshots of the Amazon EBS volumes are created as a backup
C. Ensure that the Amazon EBS volume is encrypted
D. Ensure that the Amazon EBS volumes have been pre-warmed by reading all the blocks before the test

Answer: D

Explanation: During the AMI creation process, the EC2 instance creates the snapshot of instance’s root volume and any other EBS attached to the instance. The new volumes receive maximum performance and do not

require initialization also known as pre-warming. The storage blocks from the volume, restored from snapshots must be initialized before the accessing of the block. This process takes time but increases the latency of an I/O operation for every time the first block is accessed.

34. As the primary monitoring system for your web application, you use Amazon CloudWatch. After a recent software release, when using the web application, the users get sporadic 500 Internal Server Errors. You want to create an alert CloudWatch and warn an on-call engineer when this happens. How can you do this by using AWS? (Choose 3)

A. By installing a CloudWatch logs agent on your server to stream web application logs to CloudWatch

B. By deploying your web application as an AWS Elastic Beanstalk application and using the default Elastic Beanstalk CloudWatch metrics to capture 500 internal servers then setting a CloudWatch alarm on that metric
C. By using Amazon Simple Email Service to notify an on-call engineer when a CloudWatch alarm is triggered
D. By creating a CloudWatch logs group and metric filters that capture 500 internal server errors then setting a CloudWatch alarm on that metric
E. By using Amazon simple notification service to notify an on-call engineer when a CloudWatch alarm is triggered

Answer: A, D, and E
Explanation: CloudWatch Logs is used to monitor the applications and systems via log data. It takes logs from log data; therefore, no code changes are required. Amazon CloudWatch will send an email to you via Amazon SNS. First, build and subscribed to an SNS topic. You may add this SNS topic when the CloudWatch alarm is being generated to send an email update as the alarm changes.

35. An AWS proof-of-concept feature needs to be developed quickly in a development team. Several team members are responsible for managing, supporting or only viewing the project

over time. Each user has their own IAM user account already. Instead of using individual IAM accounts, it is best if the project has a dashboard with an overall view of the project, including designing, reviewing and deploying. The team leader has asked Alas about the approaches the team should use. Which AWS system should Alas suggest?

A. AWS Elastic Beanstalk

B. AWS CodeStar
C. AWS CodePipeline
D. AWS CloudFormation

Answer: B

Explanation: For the development of the project, the AWS CodeStar project integrates other AWS products. The toolchain may include source control, build, deployment and so on depending on the AWS CodeStar project model used.

36. In order for live debugging in a highly safe environment, Richard needs account-level access to production instances. What should he do?

A. Place the credentials provided by Amazon EC2 onto an MFA encrypted USB drive, and physically share it with each developer so that the private key never leaves the office

B. Create a user account for each user on all instances and place the user's keys in the credentials file in the appropriate account
C. Place an internally created private key into a secure S3 bucket with server-side encryption using customer keys and configuration management, create a service account on all the instances using this private key, and assign IAM users to each developer so they can download the file
D. Place the credentials provided by Amazon Elastic Compute Cloud (EC2) into a secure Amazon Simple Storage Service (S3) bucket with encryption enabled. Assign AWS Identity and Access Management (IAM) users to each developer so they can

download the credential file

Answer: B

Explanation: The file is located at "./aws/credentials" on Linux, MacOS or Unix or at "C:\Users\USERNAME\aws\credntials" on Windows. An instance profile is a container for an IAM role that you can use to pass role information to an EC2 instance when the instance starts. A private S3 bucket can be created for each developer, the keys can be stored in the bucket and then assigned to the instance profile.

37. If an Auto-scaling group runs in Amazon EC2, then it easily scales up and down to load in a 10-minute window; but after load peaks, you start having issues with your configuration management system where previously finished Amazon EC2 assets are still active. What configuration management framework can you use for reliable and efficient control of the cleanup of Amazon EC2 resources? (Choose 2)

A. Use Amazon Simple Workflow Service (SWF) to maintain an Amazon DynamoDB database that contains a whitelist of instances that have been previously launched, and allow the Amazon SWF worker to remove information from the configuration management system

B. Write a small script that is run during the Amazon EC2 instance shutdown to de-register the resource from the configuration management system
C. Create an Auto-scaling group lifecycle hook to hold the instance in a terminating: wait state until your termination is complete. Once termination is complete, notify Auto-scaling to complete the lifecycle hook and move the instance into a terminating: proceed state
D. Configure an Amazon Simple Queue Service (SQS) queue for Auto-scaling actions that has a script that listens for new messages and removes terminated instances from the configuration management system
E. Write a script that is run by a daily cron job on an Amazon EC2 instance and that executes API Describe calls of the EC2

Auto-scaling group and removes terminated instances from the configuration management system

Answer: B and C

Explanation: On scaling down, if the resources are not getting terminated, you can use Lifecycle Hooks. It allows you to carry out individual actions by stopping instances when Auto-scaling Group starts or terminates them. If an instance has been terminated, it stays in a state of waiting until you either complete the life cycle action via the entire life cycle action command CLI or CompleteLifecycleAction API, or the timeout period ends.

In order to clean up all resources during this termination process, you need to execute the script and in this way, also the extra cost is not required.

Write a script and run the script at shutdown to clean resources. This means that we need to add a script inside the EC2 case and create a task scheduler job and configure it when shut-down takes place.

38. Paul has a large number of web servers in an Auto-scaling group behind a load balancer. For every visitor, he wants to collect data from the logs after processing. These logs are filtered and process on an hourly basis. After collecting data, he put the data back in a durable store to run the report. Web Servers are constantly launching and terminating according to the defined policy, and Paul does not want to lose any of the log data during this launching and termination. Which approaches can meet the demand? (Choose 2)

A. On the web server, create a scheduled task that executes a script that rotates and transmits the logs to Amazon Glacier. Ensure that the operating system shut down procedure triggers a log transmission when the EC2 instance is stopped/terminated. Use the Amazon data pipeline to process the data in Amazon Glacier and run reports every hour

B. On the web servers, create a scheduled task that executes a script to rotate and transmit the logs to an Amazon S3 bucket. Ensure that the operating system shut down procedure triggers a log transmission when the EC2 instance is stopped/terminated. Use the AWS data pipeline to move log data from the Amazon S3 bucket to Amazon Redshift in order

to process and run reports every hour

C. Install an Amazon CloudWatch Logs agent on every web server during the bootstrap process. Create a CloudWatch Log Group and define a metric filter to create custom metrics that track unique visitors for the streaming web server logs. Create a scheduled task on an Amazon EC2 instance that runs every hour to generate a new report based on the CloudWatch custom metrics

D. Install an AWS Data pipeline logs agent on every web server during the bootstrap process. Create a log group object in the AWS Data pipeline, and define metric filters to move processed log data directly from the web servers to Amazon Redshift and run reports every hour

Answer: B and C

Explanation: You can download and configure the CloudWatch Logs agent in an existing EC2 instance. You can publish your own custom metrics to CloudWatch with a CLI or an API.

Amazon Redshift is a simple and cost-effective data store for the review of all your data using the standard SQL and Business Intelligence (BI) software. You can run complicated analytical queries on petabytes of structured data using sophisticated database enhancement, high-performance column space on local drives, and massively parallel query performance. The bulk of results return in seconds. It allows you to run complex analytic queries against petabytes of structured data, using sophisticated query optimization, columnar storage on high-performance local disks, and massively parallel query execution. Most results come back in seconds.

39. A financial sector creates a range of web applications with multiple platforms and programming languages. To meet their business requirements, all applications must be developed and deployed rapidly and should be highly available. What methods should they use to quickly deploy these applications?

A. Use the AWS CloudFormation docker import service to build and deploy the applications with high availability and multiple availability zones

B. Develop each application's code in DynamoDB, and then use hooks to deploy it to Elastic Beanstalk environments with Auto-scaling and elastic load balancing
C. Store each application's code in a Git repository, develop custom package repository managers for each application's dependencies, and deploy to AWS OpsWorks in multi availability zones
D. Develop the applications in Docker containers, and then deploy them to Elastic Beanstalk environments with Auto-scaling and elastic load balancing

Answer: D

Explanation: Elastic Beanstalk supports deployment from Docker containers. Docker containers have their run time environment; you can choose your platform, programming language and application dependencies that are not supported by other platforms. By using Docker with Elastic Beanstalk, you are able to handle the details of provisioning of capacity, load balancing, scaling and health monitoring automatically.

40. There is a set of EC2 instances hosted on AWS. You create a role and assigned that role to a policy, but you are unable to use that role with any instance. What is the reason?

A. You are not able to associate an IAM role with an instance

B. You will not be able to use that role with an instance unless you also create a user and associate it with that specific role
C. You need to create an instance profile and associate it with that specific role
D. You will not be able to use that role with an instance unless you also create a user group and associate it with that specific role

Answer: C

Explanation: An instance profile is like a container for an IAM role, which is used to pass the role information when the instance starts.

41. Alex has to use AWS CloudFormation for the deployment of

the application, instead of OpsWorks and the Elastic Beanstalk. But there are some types of resources that are not supported by CloudFormation. What should he do now?

A. Use a configuration management tool such as Chef, Puppet, or Ansible

B. Specify the custom resource by separating the template into multiple templates by using nested stacks
C. Create a custom resource type using template developer
D. Specify more mappings and separate the template into multiple templates by using nested stacks

Answer: C

Explanation: Custom resources allow you to write customized provisioning logic to templates, which AWS CloudFormation runs whenever you build, modify or remove the stacks. You can use Custom resources in order to include those resources, which are not supported by AWS CloudFormation. This way, all of your associated resources can be handled in one stack.

42. A company application is running on instances with the Auto-scaling group. The instances are dynamically booted and the bootstrapping takes more than 15 minutes. You note that Auto-scaling records instances as in operation prior to the completion of the bootstrapping phase. Once you finish bootstrapping, you receive software warnings relating to new instances that cause confusion. You find the cause: your application monitoring tool is polling the Auto-scaling Service API for instances that are 'In Service', and creating alarms for new previously unknown instances. Which of the following will ensure that new instances are not added to your application monitoring tool before bootstrapping is completed?

A. Increase the desired number of instances in your Auto-scaling group configuration to reduce the time it takes to bootstrap future instances

B. Create an Auto-scaling group lifecycle hook to hold the instance in a pending: wait state until your bootstrapping is

complete. Once bootstrapping is complete, notify Auto-scaling to complete the lifecycle hook and move the instance into a pending:proceed state

C. Use the default Amazon CloudWatch application metrics to monitor your application's health. Configure an Amazon SNS topic to send these CloudWatch alarms to the correct recipients

D. Tag all instances on launch to identify that they are in a pending state. Change your application monitoring tool to look for this tag before adding new instances, and the use the Amazon API to set the instance state to 'pending' until bootstrapping is complete

Answer: B

Explanation: Lifecycle hooks allow you to carry out individual actions by stopping instances when Auto-scaling Group starts or terminates them. After adding it to Auto-scaling group, it starts working in the following ways:

- Auto-scaling responds to events by scale out or in, by launching instances and by terminating instances respectively
- Auto-scaling puts the instance into a wait state (Pending:Wait orTerminating:Wait). The instance remains in this state until either you tell Auto-scaling to continue or the timeout period ends

43. Thomas wants to deploy his multi-tier web-application by using Blue/Green deployment. Each of them has its individual infrastructure: Amazon Elastic Compute Cloud (EC2) front-end servers, Amazon ElastiCache clusters, Amazon Simple Queue Service (SQS) queues, and Amazon Relational Database (RDS) Instances. What service combination would allow him to distribute traffic among different versions of his application?

A. Create one Elastic Beanstalk application and all AWS resources (using configuration files inside the application source bundle) for each web application. New versions would be deployed updating the Elastic Beanstalk application version for the current Elastic Beanstalk environment

B. Using AWS CloudFormation templates, create one Elastic

Beanstalk application and all AWS resources (in the same template) for each web application. New versions would be deployed updating a parameter on the CloudFormation template and passing it to the cfn-hup helper daemon, and traffic would be balanced between them using Weighted Round Robin (WRR) records in Amazon Route53

C. Using AWS CloudFormation templates, create one Elastic Beanstalk application and all AWS resources (in the same template) for each web application. New versions would be deployed using AWS CloudFormation templates to create new Elastic Beanstalk environments, and traffic would be balanced between them using weighted Round Robin (WRR) records in Amazon Route53
D. Create one AWS Elastic Beanstalk application and all AWS resources (using configuration files inside the application source bundle) for each web application. New versions would be deployed using Elastic Beanstalk environments and using the Swap URLs feature

Answer: C

Explanation: For Blue/Green deployment, use CloudFormation templates with all resources for each web application. Create one Elastic Beanstalk application. Now, with the use of templates, you can deploy new versions of Elastic Beanstalk environments. For distributing traffic among different versions, you can use Route53 Weighted Round Robin. A weighted distribution enables a canary analysis to be performed where a small percentage of production traffic is entered in a new environment. The new code can be tested and errors monitored to limit the burst if problems arise. It can also scale the green environment to support the entire production load if you use elastic load balancing.

44. You created the Lambda function using a template for AWS Serverless Application Model (AWS SAM). With AWS SAM integrated into AWS CodeDeploy, you would like to move users to a new function with secure Lambda deployments. A template piece is as follows:

```
Resources:
```

```
TheLambdaFunction:
Type: AWS::Serverless::Function
Properties:
Handler: index.handler
Runtime: nodejs8.10
CodeUri: s3://mybucket/mycode.zip
AutoPublishAlias: live
DeploymentPreference:
Type: Canary10Percent5Minutes
Alarms:
# Alarms to monitor
- !Ref AliasErrorMetricAlarm
- !Ref LatestVersionErrorMetricAlarm
Hooks:
# PreTraffic: validation function before shifting the traffic
#PostTraffic: validation function after shifting the traffic
PreTraffic: !Ref PreTrafficLambdaForValidation
PostTraffic: !Ref PostTrafficLambdaForValidation
```

Which statement will prove to be WRONG after this template is used?

A. If PostTrafficLambdaForValidation has failed, the deployment is rolled back

B. If PreTrafficLambdaForValidation has failed, the deployment is rolled back
C. If an alarm such as AliasErrorMetricAlarm appears, CodeDeploy rolls back the deployment
D. Every 5 minutes, 10 percent of the traffic is shifted to the new version until all traffic has been shifted

Answer: D

Explanation: SAM is using for deploying the serverless application, integrated with CodeDeploy for save deployment. CodeDeploy supports three types of deployment Canary, Linear, and All-at-Once for Lambda

deployment.
So the statement: Canary10Percent5Minutes is wrong. This means that 10% of the entire traffic is moved immediately to the new software version. After five minutes, all other traffic is shifted.

45. An enterprise wants to migrate its application from EC2 instances to API Gateway/Lambda Serverless service. For EC2 deployment, they use Blue/Green Deployment while for the serverless application, they prefer to use canary deployment, especially for API Gateway. For example, only 10% of traffic for a new release is initially allocated to the Canary stage. How would you set up a canary API Gateway deployment?

A. Configure a canary strategy in the relevant API Gateway stage. For example, allocate 10% of the traffic to the new version. When a new API is deployed, select the stage with the enabled Canary

B. In AWS Lambda console, when a new version is published, create a Canary strategy to allocate 10% of the traffic to the new version. Add Canary stage variables if needed
C. Use Elastic Beanstalk to deploy API Gateway/Lambda services as Elastic Beanstalk has supported a Canary strategy when a new version is released for API Gateway/Lambda
D. There is no straightforward way for API Gateway to implement Canary deployment. Another service such as CodeDeploy, is needed.

Answer: A
Explanation: Canary deployment is supported by AWS API Gateway. The API is broken up into a Production release and a Canary release with a user-configured ratio in the case of a Canary release implementation. Canary settings can be located on the Stage within the API Gateway. The request distribution of the Stage can, therefore, be adjusted.

46. What should be done to store credentials on Amazon EC2 instances for connecting to an Amazon RDS MYSQL database instance?

A. Give the Amazon EC2 instance an IAM role that allows read access to a private Amazon S3 bucket. Store a file with database credentials in the Amazon S3 bucket. Have your configuration management system pull the file from the bucket when it is needed

B. Launch an Amazon EC2 instance and use the configuration management system to bootstrap the instance with the Amazon RDS DB credentials. Create an AMI from this instance
C. Assign an IAM role to your Amazon EC2 instance, and use this IAM role to access the Amazon RDS DB from your Amazon EC2 instances
D. Store the Amazon RDS DB credentials in Amazon EC2 user data. Import the credential into the instance on boot

Answer: B

Explanation: To connect the DB instance or DB cluster with IAM users or roles, am IAM policy must be created. Then attach that policy to the user or role. You can assign access to users, programs or services that usually do not have access to your AWS assets using roles. For instance, in your AWS account, you may want to allow users access to resources they do not usually have, or give users access to resources in another account on another AWS account.

47. John is using Docker to get high consistency between staging and production for the application in EC2 instance, but you are asked for de-risk deployment due to accidental inconsistencies between staging and production, which can sometimes lead to unexpected production behaviors even when stage test is passed. How do you further de-risk the rest of the execution environment knowing that AWS contains many service components that can be used beyond EC2?

A. Use AWS Config to force the staging and production stacks to have configuration parity. Any differences will be detected for you, so that you are aware of risks

B. Use AMIs to ensure the whole machine, including the kernel of

the virtual machines, is consistent, since Docker uses Linux Container (LXC) technology, and we need to make sure the container environment is consistent

C. Use AWS ECS and Dockers clustering. This will make sure that the AMIs and machine sizes are the same across both the environments
D. Develop models of your entire cloud system in CloudFormation. Use this model in staging and production to achieve greater parity

Answer: D

Explanation: After resources and stack set up, the templates can be reused to replicate the infrastructure in multiple environments.

Use the parameters, mappings and conditions parts to make templates reusable so that you can personalize the stacks when creating them. For example, you may choose a less cost-effective instance type for your development environments, in comparison to your production environments, but all other setups are the same.

48. When not running in production during all template launches, you have to automatically create a Route53 record in CloudFormation. What should you do to implement this?

A. Create two templates, one with the Route53 record value and one with the null value for the record. Use the one without it when deploying to production

B. Use a parameter for the environment, and add a condition on the Route53 resource in the template to create the record only when the environment is not production
C. Use a parameter for the environment, and add a condition on the Route53 resource in the template to create the record with the null string when the environment is production
D. Create two templates, one with the Route53 record and one without it. Use the one without it when deploying to production

Answer: B

Explanation: The optional ‘Conditions’ section contains statements that are

defined when creating a resource or defining a property. For instance, you can compare if a value is equal to a value. You can create resources based on the result of its condition. You may use conditions for reusing a model that can produce resources in a variety of contexts such as a test environment versus a production environment if you have multiple conditions. An EnvironmentType Parameter can be added to your template to accept either the prod or test as input.
You may include Amazon EC2 in the production environment with certain capabilities; however, you should use reduced capabilities to save money for the test environment. You may describe the resource and how they are created for each type of environment under conditions.

49. A finance sector has assigned a task to Thomas of implementing a Red / Black Deployment Strategy for a new EC2 project. Two similar environments have been created in AWS. The ELB with Auto-scaling Group is attached to each environment. One environment will run the live traffic and on the other, new software release will be installed. Traffic between two environments is swapped by upgrading Route53 DNS records. Which of the Red / Black deployment description from the following is not true?

A. There is no down time as the old environment can still serve traffic when the new environment is not ready

B. With Amazon Route53, only one of the versions gets traffic at any point in time
C. Route53 can switch the DNS record to both the Red and the Black environments at the same time
D. If issues arise during the deployment, rollback can be easily achieved by modifying Route53 to shift traffic back to the original environment

Answer: C

Explanation: Red/Black deployment is a newer term for Blue/Green deployment by Netflix. In that, the DNS switch can only happen to one environment at a time so the statement that it can switch to both environments is wrong.

50. Any changes to the AWS resource configurations must be traced and documented to meet company regulatory requirements. What steps are best practices for implementing Infrastructure as a Code and monitoring the health of the environment of the organization? (Choose 2)

A. In AWS System Manager “Run Command”, create a JSON command document in order to configure a customized logging system for EC2 instances

B. In AWS console, activate the versioning for S3 buckets so that every object stored in the S3 buckets is version controlled
C. Create VPC peering between two VPCs so that instances in either VPC can communicate with each other privately
D. For disaster recovery purposes, a CloudFormation template is used to build a replica of a production environment in another AWS region
E. Use AWS CLI to configure a VPC, which contains public subnets, private subnets and a bastion host, which is used for remote access to the instances in private subnets

Answer: A and D

Explanation: For resource provisioning, the best practice is to use AWS CloudFormation. In order to create the same resources whenever needed, you can use the same template. The JSON command document contains Systems Manager's parameters and action. Itcan also be re-used and version control via Git.

51. When the status of Trusted Advisor checks are changing, you are to provide a solution to notify the team properly. For example, when Trusted Advisor's cost optimization checks have just identified an Amazon EC2 instance with a low use system, the operating team should be informed by the Slack Channel to respond to the changes in status and potentially reduce cost. In order to satisfy this requirement, what two combinations should the company’s DevOps Engineer use? (Choose two)

A. Create a new SNS topic, which is in charge of providing

customized notifications to the Slack channel

B. Use a Lambda function to pass a customized notification to the Slack channel when check status in Trusted Advisor has changed
C. In AWS CloudWatch Logs, create a metric filter for any new logs, which contain “Check Item Refresh Status”
D. In the Cost Optimization dashboard of Trusted Advisor, configure a notification to an SNS topic when the status check has found a new event
E. Create a new CloudWatch Events rule. Add event source as "Trusted Advisor" and event type as "Check Item Refresh Status"

Answer: B and E

Explanation: In this scenario, a new Amazon CloudWatch Events policy is the perfect way to monitor trusted advisor’s check. And the target for this policy is the Lambda function that can design as per the requirement of when and how notifications are sent to the slack channel. Here, CloudWatch events are the best tool for monitoring the changes in Trusted Advisor.

52. Juliet’s company needs to reduce costs and all teams must develop an operational cost reduction plan, where possible. You work within the DevOps team and your team manager asked the team to think of how to optimize the cost of usage of AWS resources. Which tool from the following DOES NOT help in cost optimization?

A. AWS Trusted Advisor

B. AWS Config
C. AWS Cost and Usage Report
D. Amazon S3 Analytics

Answer: B

Explanation: While AWS Config offers an AWS resource inventory and configuration changes, it does not provide details about saving costs. It can also be used to ensure compliance, but making suggestions on how to optimize existing resources cost is not easy.

53. An organization has its application on AWS EC2 instance, and they give secure access to AWS Service APIs by defining IAM roles. Now the organization wants to fetch the API keys for using with AWS SDK. What should you as developer professional engineer do to configure the application on the instance?

A. When using AWS SDKs and Amazon EC2 roles, you do not have to explicitly retrieve API keys, because the SDK handles retrieving them from the Amazon EC2 MetaData service

B. When assigning an EC2 IAM role to your instance in the console, in the “Chosen SDK” drop-down list, select the SDK that you are using, and the instance will configure the correct SDK on launch with the API keys
C. Within your application code, configure the AWS SDK to get the API keys from environment variables, because assigning an Amazon EC2 role stores keys in environment variables on launch
D. Within your application code, make a GET request to the IAM Service API to retrieve credentials for your user

Answer: A

Explanation: When you use IAM roles, you do not need to manage security credentials for the applications access beacsue the IAM role securely allows the instance to make API calls.

54. Austin wants to deploy his new application in AWS with the Blue/Green Deployment technique. For Blue/Green Deployment he wants to use Java AWS SDK. The application requires services like ELB and Auto-scaling group for both current and new releases. But it may be possible that after the new ASG is attached to ELB, the tests fail. From the following, which is an important consideration in failed condition?

A. Raise a CloudWatch alarm to alert the team when there is a failure. Also, create a CloudWatch dashboard based on the Auto-scaling group metrics

B. Rollback the system to the original state by detaching the new ASG from ELB
C. In your Java file, collect logs and send them to AWS CloudWatch Logs, which can help with troubleshooting on the failure
D. Use SNS to send a notification to your team so that the team can react to the failure in time

Answer: B

Explanation: Any script or program that manipulates AWS resources should be auto-healed, and if something goes wrong, the device can be returned to its original state. In other words, the code must be sufficiently robust to ensure that when exceptions happen, the system can roll back automatically.

55. A company has video games, and for that, they are creating new APIs. The number of reads is 100 times more than writes. The top 1% of scores are 100 times more frequently hit than the rest of the scores. How will you as developer professional engineer design the use of DynamoDB?

A. DynamoDB table with roughly equal read and write throughput, with CloudFront caching

B. DynamoDB table with roughly equal read and write throughput, with ElastiCache caching
C. DynamoDB table with 100x higher read than write throughput, with CloudFront caching
D. DynamoDB table with 100x higher read than write throughput, with ElastiCache caching

Answer: B

Explanation: With caching, you can miss a roughly equal number of read to write beacuase the majority will hit 1% of scores. We must use AWS ElastiCache as we know that the value needs to be set to equal while using caching because it is able to cache DynamoDB query rather than a distributed proxy cache for content delivery; CloudFront cannot directly cache DynamoDB queries.

56. A large number of Lambda functions with an AWS Code Deploy service are deployed by your DevOps team. In CodeDeploy, the team has already created new apps. References to the deployment groups that include configurations used during deployment are then set up. Which methods are supported for the deployment configurations? (Choose 3)

A. CodeDeployDefault.LambdaCanary10Percent5Minutes (Shifts 10 percent of traffic in the first increment. The remaining 90 percent is deployed five minutes later)

B. CodeDeployDefault.AllAtATime (Deploys the new version to all Lambda functions at a time)
C. CodeDeployDefault.LambdaAllAtOnce (Shifts all traffic to the new Lambda functions at once)
D. CodeDeployDefault.LambdaHalf (Shifts half traffic to the new Lambda functions and then shifts the other half)
E. CodeDeployDefault.LambdaLinear10PercentEvery10Minutes (Shifts 10 percent of traffic every 10 minutes)
F. CodeDeployDefault.Immutable (Shifts the traffic to the new Lambda functions by performing an immutable update)

Answer: A, C, and E

Explanation: When the AWS CodeDeploy deployment group is created, users must adjust the rules for determining how quickly an application will be deployed which is also referred to as the "Deployment configuration".

Options A, C, and E are correct because Canary is supported for shifting 10% of traffic at first then the remaining will be shifted, Lambda AttatOnce option is also best as it shifts all traffic at a time. With Linear, you can shift the traffic between each increase in equal amounts with an equal number of minutes. You can choose from linear options that specify how much traffic has shifted in every increase and how many minutes each increase takes.

57. It takes more than four hours for your current log analysis application to report the top 10 users of your web application. You were asked to implement a system that would allow you to report this information in real time, make sure that the report is always up to date and that the number of requests for your web application has

increased so it will be able to handle it. Choose the cost-effective option that can meet the requirements.

A. Configure an Auto-scaling group to increase the size of your Amazon EMR cluster

B. Post your log data to an Amazon Kinesis data stream, and subscribe to your log-processing application so that it is configured to process your logging data, configure your application to Auto-scale to handle the load on demand
C. Publish your data to CloudWatch Logs, and configure your application to Auto-scale to handle the load on demand
D. Publish your log data to an Amazon S3 bucket. Use AWS CloudFormation to create an Auto-scaling group to scale your post-processing application, which is configured to pull down your log files stored in Amazon S3

Answer: B

Explanation: For rapid data intake and real time processing of data like logs, market data, website clickstreams, etc., you can use Kinesis Streams.

Amazon Kinesis makes collecting, processing, and analyzing data in real time easier and enables you to get timely insights and respond quickly to new information. Amazon Kinesis offers key capabilities to process streaming data in a cost-effective manner, and flexibility to choose the tools that best fit your application. With Amazon Kinesis, you can enter into real times data such as application logs, website clickstreams, IoT telemetry data, and more into your databases, data lakes and data warehouses, or build your own real-time applications using this data. Instead of having to wait for all your data to be collected before the processing can start, Amazon Kinesis can process and analyze data as it arrives and responds in real time. For handling of an increase in the load, you can use Auto-scaling for your application.

58. Management observed the increase in billing cost after reviewing the last quarter monthly bills from Amazon. On researching for the increase in cost, you found that one of the new services is doing a lot of GET bucket API calls to S3 to build a metadata cache of all objects in the application bucket. You need to provide a solution, which reduces the amount of these GET bucket

API calls. What would be your strategy?

A. Using Amazon SNS, create a notification on any Amazon S3 objects that automatically updates a new DynamoDB table to store all metadata of the new object. Subscribe the application to the Amazon SNS topic to update its internal Amazon S3 object metadata cache from the DynamoDB table
B. Update your Amazon S3 bucket's lifecycle to automatically push a list of objects to a new bucket, and use this list to view objects associated with the applications bucket
C. Create a new DynamoDB table. Use the new DynamoDB table to store all metadata of all objects uploaded to Amazon S3. Anytime a new object is uploaded, update the application's internal Amazon S3 object metadata cache from DynamoDB
D. Upload all files to an ElastiCache file cache server. Update your application to now read all files metadata from the ElastiCache file cache server, and configure the ElastiCache policies to push all files to Amazon S3 for long-term storage

Answer: A

Explanation: The best option is to have a notification, which then triggers an update to the application to update the DynamoDB accordingly. It is the best way to reduce the usage of GET requests.

59. You have a complex system, involving multiple three-tier applications, networking, and IAM policies. The requirements for the new system remain so you still do not know how many AWS components there will be in the final design. The AWS CloudFormation is used to define these AWS resources so that your infrastructure can be automated and version controlled. How will you use AWS CloudFormation to provide your customers with agile new environments in an economical and reliable way?

A. Manually create one template to encompass all the resources that you need for the system, so you only have a single template to version-control

B. Create multiple separate templates for each logical part of the

system, and provide output from one to the next using an Amazon EC2 instance running the SDK for finer granularity of control

C. Manually construct the networking layer using Amazon VPC because this does not change often and then use AWS CloudFormation to define all other ephemeral resources
D. Create multiple separate templates for each logical part of the system, create a nested stack in AWS CloudFormation, and maintain several templates to version-control

Answer: D

Explanation: As your infrastructure grows, common patterns may arise in which each template you declare has the same components. You can separate and create special templates for these common components. This allows you to mix and match various templates, however using nesting stacks to create a single, unified stack. Nested stacks are stacks that create other stacks. To create nested stacks, use the AWS::CloudFormation::Stack resource in your template to reference other templates.

60. William's team has several pipelines for AWS CodePipeline service but, even when some pipelines have failed, you have found your team is unable to get any notification. You want to set up a feature in order to receive an email and an SMS from your on-call team if the whole or a certain phase of the pipeline changes its status to FAILED. To enforce this, what combinations of services should you use? (Choose 2)

A. AWS CloudWatch Alarms

B. AWS SQS
C. AWS CloudTrail
D. AWS SNS
E. AWS CloudWatch Events

Answer: D and E

Explanation: Amazon CloudWatch Events may be used to identify and respond to the pipeline, stage or action changes. CloudWatch Events instead, based on rules, invokes for one or more target actions if a pipeline, phase, or

action enters the specified state. Then CloudWatch Events rule can be configured as a target for SNS-topic. Therefore, users will be able to create a subscription of email or SMS alerts for this new issue, so as to alert the team when the state changing is failing.

61. HTTP health check has been enabled for Elastic Load Balancing. You see that all instances take health checks after looking at the AWS management console, but your customers are informed that your site is not responding. What is the reason for this?

A. Latency in DNS resolution is interfering with Amazon EC2 metadata retrieval

B. The application is returning a positive health check too quickly for the AWS Management Console to respond
C. The health check in place is not sufficiently evaluating the application function
D. The HTTP health checking system is misreporting due to latency in inter-instance metadata synchronization

Answer: C

Explanation: The custom health check is used to evaluate the functionality of the application. If the application functionality is not working and you do not have custom health checks, the instances will still be seemed as healthy. You can send the data from your health checks to Auto-scaling if you have custom health checks, so Auto-scaling can use this data. For example, you can set the health status of the instance to "unhealthy" if you determine that the situation is not working as expected. The next time Auto-scaling conducts an instance health check, the instance will be unhealthy and a replacement instance will be launched.

62. You decided that deployments in blue/green technique would prove beneficial for your company, after a regular review with the development team. What should you do to incorporate this technique?

A. Using an AWS CloudFormation template, re-deploy your application behind a load balancer, launch a new AWS

CloudFormation stack during each deployment, update your load balancer to send half your traffic to the new stack while you test, update the load balancer to send 100% of the traffic to the new stack after verification, and then terminate the old stack

B. Create a new Auto-scaling group with the new launch configuration and desired capacity the same as that of the initial Auto-scaling group and associate it with the same load balancer. Once the new Auto-scaling group instances is registered with ELB, modify the desired capacity of the initial Auto-scaling group to zero and gradually delete the old Auto-scaling group
C. Re-deploy your application on AWS Elastic Beanstalk, and take advantage of Elastic Beanstalk deployment types
D. Using an AWS OpsWorks stack, re-deploy your application behind an elastic load balancing load balancer and take advantage of OpsWorks stack versioning, create a new version of your application during deployment, tell OpsWorks to launch the new configuration behind your load balancer, and when the new version is launched, terminate the old OpsWorks stack

Answer: A

Explanation: The blue group is used to carry the production load, and the green group is used for stage and deploy with the new code**.** When it is time to deploy, you must attach the green group to the load balancer to introduce the traffic to the new environment. The load balancer favors the green Auto-scaling group for HTTP / HTTPS listeners because it uses a less outstanding algorithm for routing requests.

63. You have been assigned to use AWS OpsWorks to implement the scalable distributed system. You have to scale up your distributed system on demand. Every node must have a configuration file that includes the hostnames of the other instances in the layer, as it is distributed. How should AWS OpsWorks be configured to manage to scale dynamically?

A. Update this configuration file by writing a script to poll the AWS OpsWorks service API service for new instances. Configure your base AMI to execute this script on the operating system start-up

B. Create a chef recipe to update this configuration file, create your AWS OpsWorks stack to use custom cookbooks, and assign this recipe to execute when instances are launched
C. Create a chef recipe to update this configuration file, configure your AWS OpsWorks stack to use custom cookbooks, and assign this recipe to the configure life cycle event of the specific layer
D. Configure your AWS OpsWorks layer to use the AWS provided recipe for distributed host configuration, and configure the instance hostname and file path parameters in your recipe's settings

Answer: C

Explanation: In the AWS OpsWorks stacks lifecycle event, each set has a layer of 5 lifecycle events, and each one is associated with a set of recipes that are specific to that layer. For recipes, you must use custom cookbooks.

For this, you need to configure a lifecycle event, which occurs on all stack's instances in the following cases:

- Attaching an ELB to a layer or detaching one from the layer
- Associating an EIP with an instance or disassociating it
- Instance entering or leaving the online state

64. A company saves its code for web applications in the Git Repository, and now needs to deploy this application in AWS, which is Node.js. How will this be done? (Choose 2)

A. Create an AWS CloudFormation template, which creates an instance with the AWS::EC2::Instance resource type and an AMI with Docker pre-installed. With UserData, install Git to download the Node.js application and then set it up

B. Create a Docker file to install Node.js and gets the code from

Git. Use the Dockerfile to perform the deployment on a new AWS Elastic Beanstalk application

C. Create an AWS CloudFormation template, which creates an instance with the AWS::EC2::Container resources type. With UserData, install Git to download the Node.js application and then set it up

D. Create an Elastic Beanstalk application. Create a Docker file to install Node.js. Get the code from Git. Use the command "aws git.push" to deploy the application

Answer: A and B

Explanation: During the launch of the instance, you can automate the configuration tasks and script runs on an instance by passing user data. You can pass two types of user data to Amazon EC2: shell scripts and cloud-init directives. You can also pass this data into the launch wizard as plain text, as a file (this is useful for launching instances using the command line tools), or as base64-encoded text (for API calls). Elastic Beanstalk supports the use of Docker container web applications. You can create your own runtime environment with Docker containers. You can choose your own platform, programming language and any software dependencies that are not provided by other systems such as package managers or tools. Dockers are independent and include all the settings and software your web application needs to run. Dockers are free of charge. So using the CloudFormation template to create an instance with pre-installed Docker AMI is a good option. You can also use Docker container with Elastic Beanstalk. In a container, you define all requirements and use this container with Elastic Beanstalk to deploy an application.

65. You use CloudFormation to organize the resources of your application. During your test phase, your Amazon RDS type instance was changed and the instance was re-created, which resulted in the loss of test data. How are you to avoid this in the future?

A. Within the AWS CloudFormation parameter with which users can select the Amazon RDS instance type, set Allowed Values to only contain the current instance type

B. Update stack using ChangeSets

C. In the AWS CloudFormation Template, set the AWS::RDS::DB instances class property to be read-only
D. Use an AWS CloudFormation stack policy to deny updates to the instance
E. Subscribe to the AWS CloudFormation notification "BeforeResourceUpdate", and call CancelStackUpdate if the resources identified are the Amazon RDS instance

Answer: D

Explanation: By default, when you create a stack, all update actions are allowed. So any one can perform updates on the stack's resources. In order to avoid any accidental delete or update of any resource in the stack, you need to use Stack Policy. It is a JSON document in which you defined a specific resource, you can perform update action by explicitly defining the ALLOW permission. When you use Stack policy, it protects all resources in the stack by default. Only one stack policy can be defined per stack and it only applies during an update of stack.

66. The use of AWS CodeStar for the central control and monitoring of a new application on an AWS system is in consideration. It is a web application, which is Node.js and deployed in a Lambda function. The project was designed smoothly in 10 minutes with a template given by CodeStar. What part of the information should you view and track in the AWS CodeStar dashboard for this project?

A. The CodePipeline status including the stages of source, build and deploy

B. The application activity status provided by AWS CloudWatch
C. The commit history from CodeCommit
D. All of the above

Answer: D

Explanation: The CodeStar project in Lambda is based on the CodeCommit tools (source), CodeBuild (build), CloudFormation (Deploy) and CloudWatch (Monitor) for a Node.js application:

The information in the AWS CodeStar dashboard are:

- Commit history provided by CodeCommit, which users can use to get a rough idea about the latest code commit
- Continuous deployment status from CodePipeline, like source, build and deploy
- CloudWatch is integrated with CodeStar, which offers an application and resource monitoring solution for the company

67. You are working for a start-up that has developed a new mobile app for photo sharing. Your application has grown in popularity in recent months, resulting in a decrease in application performance due to the increase in load. Your application features a two- tier architecture consisting of an Auto-scaling PHP application tier and MySQL RDS instance initially deployed with CloudFormation. The auto-scaling group has a min value of 4 and a max value of 8. Due to the high CPU usage of the instances, the desired capacity is now 8. Once analyzed, you are confident that performance problems stem from a CPU capacity restriction, while memory use is low. You thus decide to move from M3 instances to C3 instances that are computer-optimized.

How would you deploy this change to reduce interruptions for your end-users as minimum as possible?

A. Update the launch configuration specified in the AWS CloudFormation template with the new C3 instance type. Run a stack update with the new template. Auto-scaling will then update the instances with the new instance type

B. Update the launch configuration specified in the AWS CloudFormation template with the new C3 instance type. Also, add an UpdatePolicy attribute to your Auto-scaling group that specifies an Auto-scalingRollingUpdate. Run a stack update with the new template

C. Sign into the AWS Management Console, copy the old launch configuration, and create a new launch configuration that specifies the C3 instances. Update the Auto-scaling group with the new launch configuration. Auto-scaling will then update the instance type of all running instances

D. Sign into the AWS Management Console and update the existing launch configuration with the new C3 instance type. Add an UpdatePolicy attribute to your Auto-scaling group that specifies an Auto-scalingRollingUpdate

Answer: B

Explanation: The AWS::Auto-scaling::Auto-scalingGroup resource supports an UpdatePolicy attribute, which defines how an Auto-scaling group resource is updated when an update to the CloudFormation stack occurs. A common approach is executed, which is a rolling update for updating an Auto-scaling Group by defining the Auto-scalingRollingUpdate policy. This keeps the same Auto-scaling Group and, according to the specified parameters, replaces old instances with new ones.

68. You are responsible for an insurance company for the daily operation of the online quota system used in your company to offer insurance quotes to the public. Your company wishes to use the application logs created by the system to better understand customer behavior. You have developed a log management system that meets the following requirement: All log entries must be kept by the system even during an unplanned instance failure. The consumer monitoring teams need direct access to logs from the last seven days, and eventually, all past logs must be accessible to the fraud investigation teams, but they will wait until 24 hours before the logs are usable. What steps would you take to achieve these requirements? (Choose 3)

A. Create a house keeping script that runs on T2 micro instance managed by an Auto-scaling group for high availability. The script uses the AWS API to identify any unattached Amazon EBS volumes containing log files. Your house keeping script will mount the Amazon EBS volume, upload all logs to Amazon S3, and then delete the volume

B. Configure your application to write logs to the instances ephemeral disk, because this storage is free and has good write performance. Create a script that moves the logs from the instance to Amazon S3 once an hour

C. Write a script that is configured to be executed when the instance is stopped or terminated and that will upload any remaining logs on the instance to Amazon S3
D. Create an Amazon S3 lifecycle configuration to move log files from Amazon S3 to Amazon Glacier after seven days
E. Configure your application to write logs to the instance's default Amazon EBS boot volume, because this storage already exists. Create a script that moves the logs from the instance to Amazon S3 once an hour
F. Configure your application to write logs to a separate Amazon EBS volume with the "delete or termination" field set to false. Create a script that moves the logs from the instance to Amazon S3 once an hour

Answer: A, D, and F

Explanation: The glacier is the best option because all logs must be stored indefinitely. The data can be streamed from S3 to Glacier with Lifecycle events. You can specify the lifecycle management of objects inside a bucket by configuring a lifecycle. The configuration is a set of rules that define an Amazon S3 action for a group of objects in each rule.

We can use a cost-effective EC2 instance to ensure the minimum memory requirements for the OS and script execution. In this case, a t2.micro instance can be used, taking into account the computing resource requirements of the instance and the cost factor.

The EC2 uses the EBS volumes, and the logs are stored in EBS volumes marked for non-termination. It is one of the ways to fulfill the requirement.

69. A company wants to make easier deployment and reduce the time taken by the deployment for the improvement of deployment as they deploy five times a week at max. They hired you as the DevOps engineer to create a CI pipeline that can build AMIs. How will you do this?

A. By having the CI system launch a new instance, then bootstrapping the code and dependencies on that instance, and creating an AMI using the CreateImage API call

B. By using OpsWorks to launch an EBS-backed instance, then

using a recipe to bootstrap the instance, and then having the CI system use the CreateImage API call to make an AMI from it

C. By using a dedicated EC2 instance with an EBS Volume, then downloading and configuring the code, and creating an AMI out of that
D. By uploading the code and dependencies to Amazon S3, launching an instance, downloading the package from Amazon S3, then creating the AMI with the CreateSnapshot API call

Answer: A

Explanation: As the number of calls for deployment is less, open sources like Jenkins can be used for CI-based systems. It is used as an automation server and as a simple CI. For using the CI system to launch the instance and for user data, define bootstrap code and then create AMI by using CreateImage API.

70. John has an ELB with an Auto-scaling group and he needs to phase-out all old instances and replace them with new types of instance. What steps should he take to fulfill his task? (Choose 2)

A. Use the Newest Instance to phase out all instances that use the previous configuration
B. Attach an additional Auto-scaling configuration behind the ELB and phase in newer instances while removing older instances

C. Attach an additional ELB to the Auto-scaling configuration and phase in newer instances while removing older instances
D. Use the Oldest Launch Configuration to phase out all instances that use the previous configuration

Answer: B and D

Explanation: Auto-scaling terminates the instances that have an old configuration policy while using the OldestLaunchConfiguration policy. This is helpful when you are updating a group and phasing out the instances from the previous configuration.

71. Alan has created a web app and stores it for static website

hosting in an Amazon S3 bin. This application can access the data in Amazon DynamoDB table using the AWS SDK for JavaScript in the browser. How can you ensure that API keys are kept secure to access DynamoDB data?

A. By creating an Amazon S3 role in IAM with access to the specific DynamoDB tables, and assigning it to the buckets hosting the website

B. By configuring S3 bucket tags with AWS access keys for the bucket hosting the website so that the application can query them for access
C. By configuring a web identity federation role within IAM to enable access to the correct DynamoDB resources and retrieving temporary credentials
D. By storing AWS keys in global variables within the application and configuring the application to use these credentials when making requests

Answer: C

Explanation: You do not need custom sign-in software to build or manage your own user identities with the web identity federation. You can instead sign in with an IDP — such as login to Amazon, Facebook, Google or any other OpenID Connect (OIDC) compliant IDP— and you will obtain an authentication token, then exchange that token for temporary security credentials in AWS that map to an IAM role with permissions to use the resources in your AWS account. Using an IDP helps you keep your AWS account secure because you do not have to embed and distribute long-term security credentials with your application.

72. An educational institution uses AWS to host its application. It uses CloudFormation templates and Auto-scaling. Now, the institute observes that the number of students using applications is increasing and it is facing many performance issues. If the EC2 instance type is changed to C3, the performance will improve. What is the proper way to change the instance type?

A. Update the AWS CloudFormation template that contains the

launch configuration with the new C3 instance type. Run a stack update with the updated template, and Auto-scaling will then update the instances one at a time with the new instance type

B. Update the existing launch configuration with the new C3 instance type. Add an UpdatePolicy attribute to your Auto-scaling group that specifies an Auto-scaling RollingUpdate in order to avoid downtime
C. Update the launch configuration in the AWS CloudFormation template with the new C3 instance type. Add an UpdatePolicy attribute to the Auto-scaling group that specifies an Auto-scalingRollingUpdate. Run a stack update with the updated template
D. Copy the old launch configuration, and create a new launch configuration with the C3 instances. Update the Auto-scaling group with the new launch configuration

Answer: C

Explanation: To change the instance type, you need to change the template and also the UpdatePolicy attribute of Auto-scaling in which you specify to perform updates on Auto-scaling on updating in the stack. This is done by defining Auto-scalingRollingUpdate Policy to replace the old instance with new.

73. Louis deployed his application in AWS. The application has two ASGs attach to an ELB. From these two ASGs, ASG1 is live and ASG2 is idle. The new version is deployed on ASG2 and fully reviewed once new releases are published. If no error is found, ASG2 comes alive and ASG1 gets idle. What kind of deployment / delivery method is possible to use in this case? (Choose 2)

A. Canary Deployment
B. Blue/Green Deployment
C. Rolling Deployment
D. A/B Testing
E. Red/Black Deployment

Answer: B and E

Explanation: We know that the Blue / Green implementation has two similar surroundings and can be switched to a new environment for all traffic. In this example, the Blue / Green deployment has been carried out using two ASGs. The same concept is represented both in red / black and in blue / green. The Red version (ASG1) is created live in this case. The machine guides all traffic to ASG2 when it (the Black version) is fully operational.

74. A company has an application, in which the changes occur, but if the changes fail, then they rollback the updates, but it takes 5-6 hrs for rolling back. How will you as a DevOps engineer provide a solution to reduce the rolling back duration?

A. Use OpsWorks and re-deploy using the rollback feature

B. Use S3 to store each version and then re-deploy with Elastic Beanstalk
C. Use Elastic Beanstalk and re-deploy using Application Versions
D. Use CloudFormation and update the stack with the previous template

Answer: C

Explanation: Using Elastic Beanstalk is best for development as it quickly deploys the application in the AWS Cloud and your management complexity reduces. You simply upload your application, and the details of capacity provisioning, load balance, scaling, and health controls are automatically managed by AWS Elastic Beanstalk.

75. You have just joined an IT company to take responsibility for the management of AWS resources with another DevOps engineer. After a while, you find that an Elastic Load Balancer health check timer needs to be changed. The ELB was set up a year ago with Auto-scaling settings and you can find the CloudFormation template for this setup, which is in a CodeCommit repository. During the year, such configurations can be highly altered by somebody. However, the modification details for the CloudFormation stack are not recorded. What are you supposed to do to update the ELB health check timer?

A. Edit the template properly with the new ELB health check timer. Create a Change Set for the CloudFormation stack using the new template. If the Change Set is ok, execute the Change Set

B. Use “Detect drift” to understand the changes since the stack was created. Update the template accordingly together with the new health check timer and submit the new code to CodeCommit. Update the stack using the new template

C. Modify the template with the new Elastic Load Balancer health check timer. Update the stack using the new template. Submit the new template to the CodeCommit repository

D. Delete the CloudFormation stack since the changes in the stack are not tracked in git. Rewrite the CloudFormation template with the new health check timer and create a new CloudFormation stack using the new template. Commit the code changes to CodeCommit repository.

Answer: B

Explanation: 'Detect drift' is useful to understand the changes to the stack resources, it generally detects the changing in stack’s configuration outside the AWS CloudFormation. The CodeCommit file, which is a best practice for Infrastructure as a Code, then tracks and monitors all changes. Detect drift takes time depending on the resources in the stack.

76. You will be working for a SaaS organization as the new head of operations. Your CTO asked you to simplify and speed up debugging every part of the operation. She complains she does not know what is happening in a complex, service-based architecture because the developers only log onto the disk and with so many services, it is very difficult to find mistakes in logs. How can you best fulfill your CTO’s requirement?

A. Begin using CloudWatch logs on every service. Stream all log groups into an AWS elastic search service domain running Kibana 4 and perform log analysis on a search cluster

B. Copy all log files into AWS S3 using a cron job on each

instance. Use an S3 notification configuration on the put bucket event and publish an event to AWS Lambda. Use the Lambda to analyze logs as soon as they come in and flag issues

C. Begin using CloudWatch logs on every service. Stream all log groups into S3 objects. Use AWS EMR cluster jobs to perform adhoc map reduce analysis and write new queries when needed
D. Copy all log files into AWS S3 using a cron job on each instance. Use an S3 notification configuration on the put bucket event and publish an event to AWS Kinesis. Use apache sparks on AWS EMR to perform at-scale stream processing queries on the log chunks and flag issues.

Answer: A

Explanation: The Elasticsearch Service from Amazon makes it easy to implement, operate and scale Elasticsearch for log-analysis. The Amazon Elasticsearch Service offers the easy-to-use APIs and real-time functionality of Elasticsearch, together with the availability, scalability, and security required by production workloads. It provides integrated solutions with Kibana, Logstash, Amazon Kinesis Firehose, AWS Lambda, and Amazon CloudWatch so that you can quickly move from raw data to feasible insights.

77. In order to meet an ongoing compliance audit in the business, Wayne needs to test and set up CloudWatch alarms for the AWS resources in production like EC2, DynamoDB, Lambda, etc. To save time, he plans to use existing CloudWatch Metrics and create CloudWatch alarms accordingly. Which scenarios are suitable for this? (Choose 2)

A. When there are some critical Tomcat errors happening in EC2 instances, immediately send an SMS to the on-call support

B. When the average EC2 instance memory usage exceeds 95 percent for five minutes, trigger an alarm and launch another EC2 instance
C. When someone in the development team has provided comments on commit or pull requests in AWS CodeCommit service, notify relevant team members
D. If the latency of ELB exceeds 10 seconds over two minutes,

create an alarm and send an email notification to your team

E. Notify the team when excessive throttling occurs for a DynamoDB table, which stores users' subscription data

Answer: D and E

Explanation: As Wayne wants to use the existing CloudWatch metrics, then he should use the ELB metric, which is the Average Latency metric. On the basis of this metric, he must create an alarm and send notification via email. Then he should use the CloudWatch metric that is supported by DynamoDB for throttled read/write events or requests.

78. An organization has an application of which it released a new feature, making it a highly available application. For testing the new feature, A/B testing is used. The logs from each updated instance need to analyze real-time to observe the working. If the behavior of the log is anomalous, the instances are changes to a more stable one. What would be your strategy in this scenario?

A. Ship the logs to a large Amazon EC2 instance and analyze the logs in a live manner

B. Ship the logs to Amazon CloudWatch Logs and use Amazon EMR to analyze the logs in a batch manner each hour

C. Ship the logs to Amazon S3 for durability and use Amazon EMR to analyze the logs in a batch manner each hour

D. Ship the logs to an Amazon Kinesis Stream and have the consumers analyze the logs in a live manner

Answer: D

Explanation: For rapid data intake and real time processing of data like logs, market data, etc., you can use Kinesis Streams. You can use data collected into Kinesis Streams for simple data analysis and reporting in real time. Kinesis Streams offers a quick intake of data feed; before uploading data for intake, you do not need to batch the data on the servers.

79. Jordan needs to find a solution to cost optimization without impacting the services of AWS. What should he do? (Choose 2)

A. Turn off certain non-production instances during weekends

when no one is actually using them and turn on these instances on Monday morning

B. To save cost, use spot instances instead of on-demand instances for Jenkins, which is the production CI/CD pipeline
C. In order to decide the right size of EC2 instances, create a CloudFormation stack, with EC2 instances, to get the utilization data from CloudWatch and upload it to the Amazon Redshift cluster for right-sizing analysis. Delete the stack after the analysis result is fetched
D. For non-production environments, use default General Purpose SSD (gp2) instead of EBS Throughput Optimized HDD (st1) storage as SSD costs half the price
E. To save the cost of the VPN connections between VPC and on-premises servers, replace the VPN connections with a Direct Connect, which is cheaper and also more reliable

Answer: A and C

Explanation: Because unused EC2 instances can be stopped in order to save money, this task may automatically be executed by a script or pipeline to stop and start the instances automatically.

You can also use the Redshift cluster because this is an automated method of finding the existing data from CloudWatch and getting 3 recommendations for the right size of EC2 instances.

80. The design of a CI / CD pipeline is your responsibility for the current request. One of the main requirements is that the AWS applications should be highly available and they should auto-heal when errors occur. Which AWS resources will be helpful for you? (Choose 2)

A. Create a CloudWatch alarm based on the CPU metrics and notify the on-call team via an SNS subscription

B. Create an Elastic Load Balancer and attach an Auto-scaling group to it
C. Create an OpsWorks stack with the auto-healing feature enabled for its layer

D. Create an Elastic Beanstalk application with the auto-healing feature enabled in the environment configuration
E. Use a CloudFormation template to launch a stack for both infrastructure and application

Answer: B and C

Explanation: In order to help users to deploy fully accessible, fault tolerant and self-healing software, AWS has supplied a number of services. The key to auto-healing is to automatically start a new one whenever an error occurs when the problematic instance is stopped.

Since OpsWorks has provided its layer with an auto-healing feature, if an agent fails to communicate with the OpsWorks service for some time, AWS OpsWorks Stack automatically replaces the failed instance when the feature of auto-healing is enabled.

ELB and ASG will contribute very quickly to achieving an elastic self-healing process. If any instance does not perform well and the ELB health test fails, ASG stops and starts a new one to replace it.

81. Joel plans to use the Amazon RDS facility for fault tolerance of the application and needs its features. What is one of the benefits of Amazon RDS Multi-Availability Zone?

A. A second standby database is deployed and maintained in a different availability zone from the master, using asynchronous replication

B. A second standby database is deployed and maintained in a different region from the master, using asynchronous replication
C. A second, standby database is deployed and maintained in a different region from the master, using synchronous replication
D. A second, standby database is deployed and maintained in a different availability zone from the master, using synchronous replication

Answer: D

Explanation: The multi-AZ implementations of Amazon RDS have improved Database (DB) availability and durability, making them ideal for

work loads in production databases. Amazon RDS automatically generates a primary DB instance and synchronizes the data into a standby instance in an alternative Availability Zone (AZ) when providing the Multi-AZ DB Instance. Each AZ operates on its own physical, separate and very reliable infrastructure. An RDS would automatically fail to the standby (or read replica in Amazon Aurora), in case of an infrastructure failure, in order to restart server operations as soon as the failover is complete.

82. A company has an application in which whenever a user uploads a photo in the app, a message is sent immediately to their friends. To save the photo information per user, an AWS DynamoDB table has been created and the DynamoDB stream is enabled to capture the item changes. Which service can be integrated with DynamoDB in order to process and notify events?

A. An S3 Bucket

B. A Lambda Function
C. An SNS Notification
D. A CloudWatch Event Rule

Answer: B

Explanation: In DynamoDB, you can enable the stream feature and after that, a time-ordered sequence of item-level operations is captured, which can be accessed via its API. For Lambda function, you can configure DynamoDB Stream as a trigger as so Lambda picks up a new stream record and gets it processed.

83. The firm has a DevOps department that handles all other departments of the company's AWS accounts. A number of CloudWatch Event Rules are defined in the master AWS account to allow CloudWatch Events to be sent in other AWS accounts. In order to do this, permissions are set up in the master account, which is in region ap-southeast-2 with default CloudWatch event bus. Which entities can be added to allow submitting events on the master account in CloudWatch Event Bus? (Choose 2)

A. IAM Groups

B. An AWS Account, which creates a CloudWatch Events rule in region ap-south-1 and sends events to the master account
C. Everybody (all AWS accounts)
D. Another AWS Organization
E. Certain IAM users in the same AWS Organization

Answer: C and D

Explanation: You first need to update the privileges on the default event bus on your account to accept events from other accounts or organizations. The default event bus accepts the events from AWS services, other authorized AWS accounts, and PutEvents calls.

You may assign account IDs or Organization IDs when you change your default event bus to grant permission to other AWS accounts. Or, you can receive events of all AWS accounts. For permission, you need to select everybody, which includes all AWS accounts.

84. For the integration of thousands of Microservices, AWS ECS was used by a major IT company. There is a new management requirement that every state change must be reported to a group channel of Microsoft teams either by ECS container instances or ECS tasks. When a change occurs, the alert notifications should be readable and sent directly. Which two methods in combinations meet this requirement? (Choose 2)

A. Create a Lambda function to analyze the state change event in ECS and then provide an appropriate message to the channel in Microsoft Teams

B. Set up alarms in ECS CloudWatch Metrics. When the state of either ECS instance or ECS task changes, generate an alarm to trigger an SNS topic
C. Configure a new SNS topic called ECS_Status_Change. Register the SNS as the trigger of a Lambda function to generate custom readable alarms
D. Create a T2 micro instance to handle the customized notifications. Integrate it with Microsoft Teams via its Incoming Webhook
E. Add a new CloudWatch Events rule for all event types of ECS.

Add a Lambda function as the target when an event is triggered

Answer: A and E

Explanation: The combination of the Lambda and CloudWatch event will be the best solution for the given requirement. In this case, notifications are expected to be sent immediately when there is a state change for ECS. CloudWatch Events should be considered as it integrates with ECS and can trigger a CloudWatch Event for its target. As CloudWatch Events can trace the status change of both ECS container instance and ECS task, Lambda can be used to generate the message so you can integrate it with CloudWatch events.

85. An enterprise has a hybrid network consisting of multiple EC2 instances and Raspberry Pi devices. In order to manage all servers like patching strategy on both cloud and an on-premises environment, the enterprise plans to use AWS System Manager. The agents are also installed on all servers. Now, it wants to identify which server is the Raspberry Pi server from the Managed Instances Console. Which feature of the Raspberry Pi distinguishes it from others?

A. The machines prefixed with "i-" are Raspberry Pi devices

B. The Raspberry Pi instances have the IP addresses that are not within the VPC IP range
C. The machines prefixed with "mi-" are Raspberry Pi devices
D. The machines that have the platform name as CentOS Linux are Raspberry Pi devices

Answer: C

Explanation: Within a hybrid environment, the AWS Systems Administrator will configure EC2 instances or on-site servers. Support is provided for various Linux distributions, including Raspberry Pi devices and Microsoft Windows Server. As the "mi-" suffix is for hybrid instances, the suffix "i-" is for the Amazon EC2 instances.

86. An Auto-scaling group is configured to launch EC2 instances for the application, but you observed that the Auto-scaling group is not launching the instances in the right proportion. Instances are

launched too fast. How will you resolve this issue? (Choose 2)

A. By setting a custom metric, which monitors a key application functionality for the scale-in and scale-out process

B. By adjusting the CPU threshold set for the Auto-scaling scale-in and scale-out process
C. By adjusting the memory threshold set for the Auto-scaling scale-in and scale-out process
D. By adjusting the cool down period set for the Auto-scaling group

Answer: B and D

Explanation: The Auto-scaling cooldown period is a configuration setting that makes sure that Auto-scaling does not launch or terminate any additional instance until the previous activity takes place.

87. In a very large organization, there are multiple applications, which are constructed in different programming languages. How can you deploy these applications as fast as possible?

A. By developing each app in a separate docker container and deploying them using Elastic Beanstalk
B. By developing each app in one docker container and deploying them using Elastic Beanstalk

C. By creating a Lambda function deployment package consisting of code and any dependencies
D. By developing each app in a separate docker container and deploying them using CloudFormation

Answer: A

Explanation: EBS supports the deployment of a web application from a docker container. In docker, you can define your run time environment where you can choose your platform, programming language and application dependencies that are not supported by other platforms. Docker containers are self-contained and include all the configuration information and software your web application requires to run.

88. Aidan has a company in which he helps his IT team for setting up the AWS CodeCommit repositories. He and his team choose to pull or push code with their current SSH keys, however, two team members told him that they could NOT run successfully on their own Linux machines on CodeCommit repositories. How would he solve this problem? (Choose 2)

A. Check whether the public SSH key has been uploaded to the IAM Security Credential tab

B. If Git is used, check if the IAM users are able to access AWS CodeCommit using their AWS credentials
C. Check if the IAM user has a proper policy to access CodeCommit resource
D. In the IAM Security Credential tab, check if the user's private SSH key has been uploaded
E. Check if the IAM user has activated MFA

Answer: A and C

Explanation: CodeCommit can be accessed by several methods: HTTPS, SSH and AWS access keys. You need to ensure for SSH that IAM rules, such as AWSCodeCommitFullAccess, are associated with the IAM user. To authenticate the repository, access to the public should be uploaded.

With SSH connections, you create public and private key files on your local machine that Git and CodeCommit use for SSH authentication.

89. A number of instances are running on your OpsWork stacks. If you want to install security updates, what is the AWS recommendation in accordance with this task? (Choose 2)

A. Create a CloudFormation template, which can be used to replace the instances

B. Create and start new instances to replace your current online instances. Then delete the current instances
C. On Linux-based instances in Chef 11.10 or older stacks, run the Update Dependencies stack command
D. Create a new Opswork stack with the new instances

Answer: B and C

Explanation: By default, after an instance booting is completed, AWS OpsWorks stacks automatically install the most recent updates during setup. AWS OpsWorks stacks do not automatically install updates after an instance is online, in order to prevent interruptions such as restarting of the application server. Instead, you should manage online updates yourself so that any disturbance can be minimized.

You may use one of the following for an update:

- Create and launch new instances to replace your current online instances. Then delete the existing instances
- Run the command Update Dependencies Stack on Linux based instances in Chef 11.10 or previous stacks, which installs a current set of security patches and updates in the instance you specified

90. AWS CodeDeploy is used to manage several deployment phases, including feature testing, system integration testing, and manufacturing. Every phase has a CodeDeploy deployment group. Each group is supplemented with tags with relevant targets. Like servers containing the product, tag is deployed only if the Production Deployment Group creates a new deployment. Which targets CANNOT be added to Deployment Group by Tags in AWS CodeDeploy? (Choose 2)

A. Red Hat Enterprise Linux (RHEL) on-premises instances

B. EC2 Instance (Microsoft Windows Server 2016)
C. Lambda Function written in Python 3.6
D. EC2 Auto-scaling Groups
E. ECS Cluster with an Amazon Linux AMI

Answer: C and E

Explanation: Users can select the targets via tags during the creation of Deployment Group in CodeDeploy from any Amazon EC2 Auto-scaling Groups, Amazon EC2 instances or on-site instances but Lambda function and ECS cannot be selected.

91. The Docker Containers for a Java Application is controlled by AWS ECS within your business. Several important features have recently been developed to satisfy the market requirements and the team is under pressure to deliver new releases. When the ECS role comes into stopped status, the manager asks you to alert the team both on the slack channel and email so that the relevant staff can take immediate action. What is the best way for you to do this?

A. Configure the ECS cluster instances to send log files in /var/log/docker/ to CloudWatch Logs. Create a Metric Filter to search for the keyword “STOPPED”. Create an alarm to trigger an SNS notification whenever there is a match. Subscribe a Lambda function to the SNS topic, which sends a message to the Slack channel and email list

B. Create a CloudWatch Event rule as below. Use a Lambda function as the target to send notifications to both the Slack channel and the relevant email list.

```
{
"source": [</p>
   "aws.ecs"</p>
],
"detail-type": [
  "ECS Container Instance State Change"
]
"detail": {
  "desireStatus": [
              "STOPPED"
            ]
      }
}
```

C. Configure the ECS cluster instances to send log files in /var/log/ecs/ to CloudWatch Logs. Create a Metric Filter to search for the keyword “STOPPED”. Create an alarm to trigger

a Lambda function whenever there is a match. The Lambda function sends customised notifications to the Slack channel and email list

D. Create the following CloudWatch Event rule. Add a target of Lambda function to send a notification in Slack and another target of SNS to send the email.

```
{
"source": [</p>
    "aws.ecs"</p>
],
"detail-type": [
   "ECS Task State Change"
]
"detail": {
   "lastStatus": [
  "STOPPED"
]
        }
}
```

Answer: D

Explanation: In order to provide immediate notification, you can integrate Amazon ECS with CloudWatch events. Once ECS reaches the "STOPPED" state, the rule will immediately trigger the targets (Lambda and SNS) with this CloudWatch Event rule.

92. Phillip is using AWS Systems Manager to maintain EC2 instances. For example, for instances with a tag of "QA", you can run command to execute a shellscript. However, you want to limit the utilization of the "Run Command" feature for certain IAM users on security issues. For these specific users, you need an IAM policy to only allow them to run command for instances that have the “department” tag of “dev1” or “dev2”. Which IAM policy can help you to achieve this requirement?

A. {
"Version":"2012-10-17",
"Statement":[
{
"Effect":"Allow",
"Action":[
"ssm:SendCommand"
],
"Resource":"*",
"Condition":{
"StringLike":{
"ssm:resourceTag/department":[
"dev1"
]
"ssm:resourceTag/department":[
"dev2"
]
}
}
}
]
}

B. {
"Version":"2012-10-17",
"Statement":[
{
"Effect":"Allow",
"Action":[
"ssm:RunCommand"
],
"Resource":"*",
"Condition":{
"StringLike":{

```
                "ssm:resourceTag/department":[
                   "dev1"
                ]
                "ssm:resourceTag/department":[
                   "dev2"
]
            }
         }
      }
]
}
```

C.

```
{
 "Version":"2012-10-17",
 "Statement":[
    {
       "Effect":"Allow",
       "Action":[
          "*"
       ],
       "Resource":"*",
       "Condition":{
          "StringLike":{
             "ssm:resourceTag/department":[
                "dev1"
             ]
             "ssm:resourceTag/department":[
                "dev2"
]
          }
       }
    }
]
}
```

D. {
"Version":"2012-10-17",
"Statement":[
{
"Effect":"Allow",
"Action":[
"ssm:SendCommand"
],
"Resource":"*",
"Condition":{
"StringNotEquals":{
"ssm:resourceTag/department":[
"dev1", "dev2"
]
}
}
}
]
}

Answer: A

Explanation: If the Systems Manager Run Command needs a restriction, the best practice is to develop the appropriate IAM policy for IAM users and groups.

Option A allows only tags 1 and 2 to run the command.

93. There is a number of CloudFormation templates. Most of them are used to manage AMI IDs for various regions and instance types with mappings. The AMI IDs can, however, change regularly; for example, when there are software updates. In that case, all related CloudFormation templates must be updated, which takes time. You plan to have an automated way to search and get the correct CloudFormation AMI IDs. Which two strategies will help you to accomplish this together? (Choose 2)

A. Add a new AMI ID parameter in the CloudFormation

templates. When creating CloudFormation stacks, input the latest AMI ID parameter

B. Use a CloudWatch Event rule to execute a Lambda function every day to get the latest AMI IDs
C. In the CloudFormation template, create a custom resource type to invoke and send input values to a Lambda function in order to get the correct AMIs. After the custom resource gets a proper response, the stack proceeds with other resources
D. Prepare a shell script to fetch the latest AMIs for any region and instance type by using AWS CLI commands such as ec2 describe-images
E. Create a Lambda function to get the latest AMIs for a given region and instance type

Answer: C and E

Explanation: As they want the automated method of fetching AMI IDs for EC2 resources, creating custom resource can associate a Lambda function to get the AMI IDs. The stack can only proceed after the AMIs are received from the Lambda. This Lambda function can use region and instance type as inputs and get the related AMIs. It then returns the IDs to the custom resource.

94. A new employee has joined AWS codecommit in several AWS region for the management of code repositories. To link to SSH's CodeCommit, the newest version of a GIT has been installed in the Linux machine, an ssh-keygen key pair has been created and a public key has been added to its IAM users' security identifiers. But still, no repositories using SSH can be cloned. What ~/.ssh/ config file does the Linux computer need to create?

A. Host git-codecommit.*.amazonaws.com

User IAMUSERIDEXAMPLE
IdentityFile ~/.ssh/id_rsa.pub

B. Host git-codecommit.*.amazonaws.com

User APKAXXXXXXEXAMPLE
IdentityFile ~/.ssh/id_rsa

C. Host git-codecommit.ap-south-1.amazonaws.com

User APKAXXXXXXXEXAMPLE
IdentityFile ~/.ssh/id_rsa

D. Host *

User IAMUSERIDEXAMPLE
IdentityFile ~/.ssh/id_rsa

Answer: B

Explanation: GIT SSH Access is a common approach for connecting from a remote server to AWS CodeCommit. So, Option B is the correct configuration for SSH.

95. Arthur has a large number of AWS accounts that are managed by him. Some AWS accounts belong to AWS Organization and some do not. When there is a new EC2 termination event for any other account, it is necessary to set up the CloudWatch rule in the master account. In each account, you plan to set a CloudWatch Events rule, and send the event to the master account in the default Event Bus. What is NOT the solution to work with in this requirement?

A. In CloudWatch Events, in order to send the event to the master account, the target should be configured as “Event bus in another AWS account”

B. A proper IAM role is needed for the sender account to send events to the Event bus in the master account
C. All accounts should belong to certain AWS Organizations
D. A master account should permit the CloudWatch events to be received from other accounts

Answer: C

Explanation: The CloudWatch Event Bus is a helpful service in order to send / receive CloudWatch events to / from other AWS accounts. For that,

you need permission from the master account for events to be received.
For CloudWatch Events to submit the event, the sender must have a suitable IAM role. We also need to configure the target as an “Event bus in another AWS account” in order to send the event to another account’s default Event bus.

96. A big financial company runs a Hybrid system. Separate deployment mechanisms for the on-site and AWS servers have previously been used. The deployment management must be increased in all instances with the same method. For this necessity, you suggested using AWS CodeDeploy because all on-site machines run under Ubuntu LTS 16.04. To configure the instances on site with AWS CodeDeploy, what pre-conditions must be met? (Choose 2)

A. A VPN connection or a Direct Connect should be established between the on-premises environment and AWS VPC

B. The local account used to configure the on-premises instance must be able to run either as a sudo or root
C. The IAM identity to register the on-premises instance in CodeDeploy service must be granted proper permissions
D. The on-premises instance must open port 80 for the outbound traffic to connect to public AWS service endpoints
E. Java 8 should be installed in order for the CodeDeploy agent to work properly

Answer: B and C
Explanation: AWS CodeDeploy assists in handling on-site deployment and EC2 deployment. Nevertheless, there are certain pre-conditions for on-site servers to be met:

- For the CodeDeploy system, the administrative control is needed otherwise the deployment will not succeed
- Suitable permission must be granted to IAM Identity in order to register the instance with AWS CodeDeploy Service

97. All instances in an organization have certain latest security patches that need to be installed in time. Nevertheless, it is necessary

for patches to be installed in dev and test environment for one week before deployment in production instances. It is possible to distinguish all EC2 instances by tags. How can the AWS System Manager be implemented in the best way possible?

A. Use the Systems Manager session manager to configure the patching for EC2 instances. Apply required patches accordingly after you remotely login into the server

B. Use a pre-defined default Patch Baseline. Add tags of dev, test, and production to relevant EC2 instances. Associate the Patch Baseline with EC2 instances via tags
C. Tag the instances using dev, test, and production. In Systems Manager, run the command of “AWS-InstallPatches” based on the tags
D. Create a customized Patch Baseline. Create several Patch Groups for dev, test, and production instances. Associate the Patch Groups with the new Patch Baseline. Schedule the patching in a maintenance window as required

Answer: D

Explanation: The patch manager in the Systems Manager should be the best tool to apply patches. As some patches need to be selected, in this case, it is highly possible to first create a custom patch baseline. Create multiple patch groups for instances of dev, test, and production. Associate the new Patch Baseline with the Patch Groups. User Groups make it possible to apply the correct patches to the appropriate instances in the corresponding patch baseline rule.

98. A CloudFormation stack has included several AWS resources including EC2 instances. In the past, AMI ID management for different regions and instance types was done via a mapping table. The AMI ID is obtained in the function "Fn::FindInMap" during the creation of the EC2 resource. In this method, however, if a new AMI ID is available, the mapping table must be updated. You look for better ways to get your updated AMIs automatically. You already have a Lambda feature, which can obtain the latest AMI with the input as a region and instance type. How can the Lambda function

best be used in order to achieve this?

A. Manually trigger the Lambda function and store its output in the Systems Manager Parameter Store. Modify the CloudFormation template to get the latest AMI ID from the Parameter store

B. Create a Custom resource in the CloudFormation template. In its RequestId property, specify the name of the Lambda function to associate it with the Custom resource
C. Create a Custom resource in the CloudFormation template. Associate the Lambda with the Custom resource by specifying the Amazon Resource Name (ARN) of the Lambda function for the ServiceToken property
D. Before updating/creating the CloudFormation stack, use a shell script to run the Lambda function to get the correct AMI ID. Use the ID as a parameter for the CloudFormation template

Answer: C

Explanation: This case requires the best way to use the Lambda, meaning that several methods will work in theory. Nonetheless, we must find out which alternative requires less manual effort and is easier to implement. Lambda can provide the correct AMI ID with the Custom asset and continue creating EC2 based on it.

99. The Jenkins server was maintained by your team in an EC2 instance. The Jenkins server is used mainly to build Java-based artefacts. Currently, this Jenkins server has no Disaster Recovery Strategy. A new disaster recovery plan must be drawn up for your team, both RTO and RPO, within 24 hours. What strategy should the team choose to recover from disasters?

A. Backup & Restore strategy. For example, backup the Jenkins configuration files every day to an S3 bucket. Use CloudFormation templates to provision the necessary AWS resources

B. Hot Standby (Multi Site) strategy. Launch a fully operational

EC2 instance for the Jenkins server. Suspend its tasks unless a failover happens

C. Warm Standby strategy. Launch a smaller size EC2 instance in the same region but different VPC as a standby
D. Pilot Light strategy. Use an AMI to launch an EC2 instance in another region and stop the instance to save cost. Start the Pilot Light instance when there is an outage in the original Jenkins server

Answer: A

Explanation: The EC2 has installed a Jenkins server that is used as a CI / CD database in this situation. It is not a production application, so it has a very low impact. There is plenty of time to recover the server using a backup when there is an outage as RPO and RTO are both 24 hours. This scenario is sufficient for the Backup & Restore strategy. There are various server backup methods, such as the Jenkins configuration files, EC2-instance AMI, daily EBS snapshots, etc.

100. A Service is designed that aggregates the clickstream data in batch and deliver reports to the subscriber through the emails. The data is extremely spikey, high-scaled, geographically distributed and unpredictable. How will you design this system?

A. Use a large shift cluster to perform the analysis, and a fleet of Lambdas to perform record inserts into the Redshift tables. Lambda will scale rapidly enough for the traffic spikes

B. Use API Gateway invoking Lambdas, which Put records into Kinesis, and EMR running spark performing Get record on Kinesis to scale with spikes. Spark on EMR outputs the analysis to S3, which are sent out via email
C. Use a CloudFront distribution with access log delivery to S3. Clicks should be recorded as query string GETs to the distribution. Reports are built and sent by periodically running EMR jobs over the access logs in S3
D. Use AWS Elasticsearch service and EC2 Auto-scaling groups. The Auto-scaling groups scale based on click throughput and stream into the Elasticsearch domain, which is also scalable.

Use Kibana to generate reports periodically

Answer: C

Explanation: The ideal approach of getting the data onto EMR is to use S3. Since the data is extremely spikey and highly-scaled, using edge location through CloudFront distribution is the best way to fetch the data. When you are building the report or analyzing data from a large data set, you need to define EMR because this service is built on the Hadoop framework, which is used to process a large set of data.

101. Eugene had his application migrate to AWS. This app is in the ap-southeast-1 region. The DevOps team used a warm-standby approach in the disaster recovery strategy and built an additional functional environment in the ap-southeast-2 region. If a production system is inefficient and failure is required, what steps can be seen as effective recovery measures? (Choose 2)

A. Check Amazon Route 53 to make sure that all traffic is routed to region ap-southeast-2

B. Select tools to backup data into AWS S3. Enable encryption for sensitive data
C. Use EBS Lifecycle Manager to regularly create EBS snapshots
D. Adjust Auto-scaling groups to accommodate the increased traffic
E. Create custom AMIs and start the application in Amazon EC2 instances

Answer: A and D

Explanation: We already know that warm-standby is an extension of Pilot light. The standby is fully operational, but the system uses a minimum database size to reduce costs. You need to adjust the ASG for this requirement. The DNS records on Route53 that require manual changes or a health check is set up to automatically help with failover. This is an important step in ensuring that the site is fully restored.

102. In its AWS account and on-premises, your organization has stored a large amount of data. The data includes daily transactions,

data from clients, etc. You have a responsibility to build an AWS QuickSight service that allows other teams to analyze data, view data via dashboards, and find hidden trends through machine learning technologies. What data cannot be supplied for analysis by QuickSight?

A. A MariaDB RDS instance

B. YAML files stored in S3 buckets
C. An AWS Redshift cluster
D. A MySQL 5.1 database in the customer's data center, which is internet accessible

Answer: B

Explanation: AWS QuickSight can allow analyzing by means of machine learning techniques by using various relational data stores as a data source. These include CSV/TSV, XLSX, ELF/CLF, and JSON.

103. Russell has shifted his services to AWS. A failure of the AWS hardware that affected one of the EBS volumes was observed last week. The AWS Personal Health Dashboard warned of this problem but the team needed several hours to process this information. Now, Russell is searching for a solution from the AWS Personal Health Dashboard that notifies the group when an open issue is happening. Which solution are you going to recommend?

A. Create a new CloudWatch Event, which monitors Trusted Advisor service and trigger an SNS notification when there is a new issue

B. Configure a Lambda function, which periodically checks open issues via AWS Health API and triggers an SNS notification if a new issue has been found
C. In AWS Personal Health Dashboard console, configure an SNS notification when specific open issues appear
D. Create a CloudWatch Event, which monitors AWS Health Event and trigger an SNS notification when there is a new event

Answer: D

Explanation: The AWS Personal Health Dashboard is able to show AWS data that might be important to your AWS resources such as open issues, scheduled changes and event logs. Therefore, on the AWS Health Service basis, a CloudWatch rule is created and users can identify a particular event service, such as EBS.

104. You have customized some system and application logs in EC2 instances and delivered them to several Log Groups in CloudWatch Logs. You find that the AWS CloudWatch Logs Console just makes it very difficult to locate useful information. You prefer the logs to a downstream processing system, which can provide the operating team with more reliable and important information. To which services could CloudWatch logs be configured and data streamed? (Choose 2)

A. AWS DynamoDB

B. AWS S3
C. AWS Lambda Function
D. AWS CloudTrail
E. AWS Elasticsearch

Answer: C and E

Explanation: In order to deliver the log events to any other service, you can use Subscription in CloudWatch logs in real-time. AWS Logs has supported streaming data to Lambda that’s why you can use it or you can use AWS Elasticsearch with AWS CloudWatch logs to stream the logs data.

105. An application is deployed in an AWS Auto-scaling group in a financial company. The request traffic is most often quite smooth and the ASG is not scaled in or out. But EC2 instances may fail the Elastic Load Balancer health check, and then be terminated by ASG, as per the latest release. In order to access the instance and solve the problem, the development team has created an ASG life-cycle hook. After the instance, which AWS services can be configured to automatically receive notification in the "Terminating: Wait" state? (Choose 2)

A. AWS CloudWatch Events

B. AWS CloudWatch Alarms
C. AWS SNS

D. AWS SQS
E. AWS CloudTrail

Answer: A and C

Explanation: The auto-scaling lifecycle hook offers users the opportunity to suspend the scaling process. It is recommended that CloudWatch Event receive notifications when an event occurs and that SNS may then be added as the notification target with a the-notification-target-arn option in the AWS CLI command put-lifecycle-hook. The ARN of the target notification can be an SQS queue or SNS subject according to —notification-target-arn.

106. Paul is using the AWS platform to deploy a website using AWS ELB and ASG. The ELB is configured with the following health checks: Ping Target: HTTPS:443/healthcheck.htm, Timeout: 20 seconds, Interval: 30 seconds, Unhealthy threshold: 3 and healthy threshold: 3.

The Auto-scaling group also uses ELB as its health check type. Last weekend, there was very high traffic due to a promotional event. Nonetheless, on launching new instances by ASG, the ELB health check failed because of the traffic congestion and eventually was again terminated by ASG. As a result, it was difficult to launch new instances, which made the matter worse. Which approaches can help resolve this issue under high traffic? (Choose 2)

A. Modify Ping target to use HTTP instead of HTTPS

B. Increase the Interval to 35 seconds
C. Decrease the Unhealthy threshold to 2
D. Increase the Healthy threshold to 5
E. Increase the Timeout to 25 seconds

Answer: B and E

Explanation: In the configuration of the health checks for the Classic Load Balancer, you can set health check parameters in the console. We should provide the EC2 instance with more time to pass a health check in this scenario; The timer for ELB is regulated by Timeout to test the set Ping Target. By increasing the timer, you can allow the EC2 instance to react to the ELB health check successfully. So, Option B and E are the same but Option B gives more time to the EC2 instance for responding.

107. To perform ad-hoc business analytics queries on well-structured data, the data comes in constantly at a high velocity. Knowing that your business intelligence team understand SQL, which service(s) should you look at first?

A. Kinesis Firehose + RDS

B. EMR using Hive
C. EMR running Apache Spark
D. Kinesis Firehose + Redshift

Answer: D
Explanation: Kinesis Firehose is the easiest way to load streaming data into AWS. It can capture, transform and load the streaming data into Kinesis analytics, S3, Redshift, and Elasticsearch service. Whereas Redshift is a fully managed, petabyte scale data warehouse service in the cloud. This enables you to use your data to acquire new insights for your business and customers.

108. A company has given you the task to configure an AWS Elastic Beanstalk work tier for easy debugging but you are facing problems in finishing queue jobs. What should you do to overcome this issue?

A. Configure a Dead Letter Queue
B. Configure Enhanced Health Reporting
C. Configure Blue-Green Deployments
D. Configure Rolling Deployments

Answer: A
Explanation: Elastic Beanstalk worker environment supports Amazon SQS Dead Letter Queues. In the dead letter queue, other queues can send messages that for some reasons could not be processed. Messages that are

unsuccessful in the processing are targeted from the source queue to the dead-letter queue. You can gather these types of messages in dead-letter queues to find the reason for their failure.

109. The highest possible network performance is required for cluster computing applications. Allen already selected homogenous instance type supporting 10Gb enhanced networking. He made sure that the workload is network bound and put the instances in the placement group. What would be the last optimization that he should make?

A. Segregate the instances into different peered VPCs while keeping them all in a placement group, so each one has its own internet gateway

B. Bake an AMI for the instances and relaunch, so the instances are fresh in the placement group and do not have noisy neighbors
C. Use 9001 MTU instead of 1500 for jumbo frames, to raise packet body to packet overhead ratios
D. Turn off SYN/ACK on your TCP stack or begin using UDP for higher throughput

Answer: C

Explanation: Jumbo packet allows the data more than 1500 bytes by increasing the payload size per packets. And increasing the percentage of the packet that is not the packet overhead. The same amount of usable data requires fewer packets to send. However, you will experience a maximum trajectory of 1500 MTUs apart from a given AWS region, a single VPC or a VPC peering connection. VPN connections and traffic sent over an internet gateway are limited to 1500 MTU. If packets are over 1500 bytes, they are fragmented, or they are dropped if the Don't Fragment flag is set in the IP header.

110. Richard wants to automate the 3 layers of large cloud deployment, make it capable of tracking all the changes over time in deployment, and carefully control any alteration. Which is the best way to achieve these requirements?

A. Use OpsWorks stacks with 3 layers to model the layering in your stack

B. Use AWS config to declare a configuration set that AWS should roll out to your cloud
C. Use Elastic Beanstalk Linked Applications, passing the important DNS entries between layers using the metadata interface
D. Use CloudFormation Nested Stack templates, with three child stacks to represent the three logical layers of your cloud

Answer: D

Explanation: When your infrastructure is growing, common patterns emerge in each template, where you declare the same components. You can separately create templates for these common components. Thus, you can mix and match the various templates and also create a single unified stack using nested stacks. Stacks that are nested are stacks that create additional stacks. Use AWS::CloudFormation::Stackresource to link other templates in your template to build nested stacks.

So as the question said that Richard needs to automate the stack over a period without recreating the stack when there are any changes, then Nested stack is the best tool to reuse Common Template Pattern. AWS also recommends that nested stack updates should be executed from the parent stack.

AWS CloudFormation updates the top-level stack and initiates an update to its nesting pillars when you upgrade the top-level stack. AWS CloudFormation updates nested stack resources but does not update the resources of nested stacks that have remained unmodified.

111. In the region, ap-southeast-1 Steve has migrated a MySQL premise database to AWS RDS MySQL. Key information like customer number, address, and date of birth have been stored in the database. A read replica in the ap-southeast-2 area was created to decrease the read load in the master DB instance. It can also be supported as a data recovery scheme if the DB source instance fails. In what scenarios can the Read Replica recover data successfully? (Choose 2)

A. An AWS regional failure happening in both ap-southeast-1 and

ap-southeast-2

B. The master database instance has been mistakenly deleted
C. AWS account is compromised
D. Data has been deleted mistakenly by a bug in application code
E. An RDS hardware failure on the master instance

Answer: B and E

Explanation: Users can select another region when creating the Read Replica so that updates to the source DB instance are copied into the Read Replica in this new region. This Read Replica can be promoted as a new, standalone instance if necessary. So if RDS hardware failure occurs, it only effects master node and you can promote Read Replica to recover DB. Or, if the master database is accidentally deleted, you can use Read Replica as functional DB.

112. There is a serious outage at AWS. The EC2 is not affected, but the EC2 instance deployment script stopped working with the outage in the region. What can be the reason?

A. S3 is unavailable, so you cannot create EBS volumes from the snapshot you used to deploy new volumes
B. The AWS console is down, so your CLI commands do not work
C. AWS turns off the deploy code API call when there are major outages, to protect from system floods
D. None of the other answers make sense. If EC2 is not affected, it must be some other issues

Answer: A

Explanation: The EBS snapshots are stored in S3. If you write the script, which deploys EC2 instances, then the EBS volume needs to be constructed from snapshot stored in S3.

By using a point-in-time snapshot, you can protect information for Amazon EBS volumes in Amazon S3. Snapshots are incremental backups that only store the frames on the system that have been modified after your last snapshot. The time needed for creating the snapshot is reduced and storage

costs are saved by not duplicating the data. Only the data unique to that snapshot can be removed when you delete your snapshot. All information required to restore the data (from the moment the snapshot has been taken) to the new EBS volume, is provided in every snapshot.

113. There is an Asynchronous application using Auto-scaling and Amazon SQS. The Auto-scaling scales as per the depth of the queue. This results in the completion velocity going down, and the Auto-scaling group size to be maxed out, but there is no increase in inbound velocity. What is the reason?

A. The routing table changed, and none of the workers can process events anymore

B. Someone changed the IAM role policy on the instances in the worker group and broke permissions to access the queue
C. The scaling metric is not functioning correctly
D. Some of the new jobs coming in are malformed and unprocessed.

Answer: D

Explanation: As the velocity in not increasing then the only reason is that the new jobs coming in are malformed and unprocessed. Option D is correct because in all the other options, no job is getting completed.

114. All modifications in consumer banking information have to be audited in a file. To save this customer banking information, you need DynamoDB. Due to server failures, it is important not to lose any data. What can you do to achieve this?

A. Before writing to DynamoDB, do a pre-write acknowledgement to a disk on the application server, removing sensitive information before logging. Periodically rotate these log files into S3

B. Use a DynamoDB stream specification and periodically flush to an EC2 instance store, removing sensitive information before putting the objects. Periodically flush these batches to

S3

C. Use a DynamoDB stream specification and stream all changes to AWS Lambda. Log the changes to AWS CloudWatch logs, removing sensitive information before logging
D. Before writing to DynamoDB, do a pre-write acknowledgement to disk on the application server, removing sensitive information before logging. Periodically pipe these files into CloudWatch logs

Answer: C

Explanation: You are able to execute Lambda Functions by using DynamoDB table streams as a trigger. Triggers are customized actions in response to DynamoDB table changes. You must first activate Amazon DynamoDB Streams in your table to create a trigger. Then, to process the updates published on this stream, you must create a Lambda function.

115. How can you achieve a gigabit network throughput on EC2, when you already selected cluster-compute, 10 GB instances with enhanced networking, network-bounded workload but you do not detect a 10-gigabit speed?

A. Enable biplex networking on your servers, so packets are non-blocking in both directions, and there is no switching over-head
B. Use a placement group for your instances so the instances are physically near each other in the same availability zone
C. Ensure the instances are in different VPCs so you do not saturate the internet gateway on any one VPC
D. Select PIOPS for your drives and mount several so that you can provision sufficient disk throughput

Answer: B

Explanation: A placement group consists of a logical group of instances within the single AZ. For applications benefiting from the low network latency, high-network performance or both, placement groups are recommended. Choose an instance that supports enhanced networking to ensure the lowest latency and the maximum packet per-second network performance for your placement group.

116. Your CTO asked you to ensure that you know what all AWS account users are doing at all times to change the resources. She needs to have a report as wide as possible, on who does what over time, sent to her once a week. How are you going to do that?

A. Use CloudWatch events rules with an SNS topic subscribed to all AWS API calls. Subscribe the CTO to an email type delivery on this SNS topic

B. Use AWS IAM Credential reports delivering a CSV of all uses of IAM user tokens over time to the CTO
C. Use AWS Config with an SNS subscription on a Lambda, and insert these changes over time into a DynamoDB table. Generate reports based on the contents of this table
D. Create a global AWS CloudTrail Trail. Configure a script to aggregate the log data delivered to S3 once per week and deliver this to the CTO

Answer: D

Explanation: AWS CloudTrail is a service that helps you enabling governance, compliance and operational and risk auditing on your account. CloudTrail is used to view, search, archive, download, analyze and respond to the account activity across your AWS infrastructure. A user, role, or AWS services activities are recorded as CloudTrail events. Actions in AWS Management Console, AWS Command Line Interface and SDKs and APIs are included in events.

117. A CloudFormation template has been used by a DevOps developer to build an RDS source for a new web app. A PostgreSQL engine was used by the RDS server and encryption is disabled. Nonetheless, the database needs to be updated to enable encryption for certain security considerations. The CloudFormation template is updated accordingly (StorageEncrypted is true). The data should be restored from the latest DB snapshot in order to prevent data losses during the update of the CloudFormation stack. In accordance with that necessity, which two steps should be taken? (Choose 2)

A. Add the DBSnapshotIdentifier property with the ID of the used

DB snapshot

B. Make sure that automated snapshots are working properly and record the last snapshot ARN ID
C. Add a DeletionPolicy of Snapshot in the CloudFormation template
D. Add a Stack Policy in the CloudFormation stack to prevent the DB resource from being deleted
E. Deactivate any applications that are using the DB instance and then create a manual snapshot

Answer: A and E

Explanation: The DB instance will be deleted and replaced by a new one if StorageEncrypted is modified and the CloudFormation updated. Meanwhile, the DBSnapshotIdentifier template should be used to point to the DB snapshot used by the CloudFormation stack. As the manual snapshot is able to maintain the data well, the CloudFormation template will use the snapshot ARN. From the snapshot defined by DBSnapshotIdentificier, AWS CloudFormation can create a new database.

118. An organization has 2000-engineer and plans to use AWS for the first time on a large scale. Now, the organization wants to integrate its identity management system, which is running on Microsoft active directory with AWS. As you know that the 2000 engineers are the power-users of Active Directory, how can you easily manage AWS Directories?

A. By using AWS Directory Service Simple AD

B. By using a sync domain running on AWS Directory Service
C. By using an AWS Directory Sync Domain running on AWS Lambda
D. By using AWS Directory Service AD Connector

Answer: D

Explanation: AD Connector is a directory gateway, by which you can redirect the directory request to your existing Microsoft Active Directory without caching any information in the cloud. It is available in two sizes:

large and small. Small AD Connector is designed for the organization up to 500 users while the larger supports up to 5000 users in an organization.
With an AD connector, you can get multiple benefits like:

- Your end users and IT administrators can access AWS applications such as Amazon WorkSpaces, Amazon WorkDocs or Amazon WorkMail using their existing credentials
- You can also use it to enable multi-factor authentication, by integrating into your existing RADIUS-based MFA infrastructure, to provide a further layer of security in access to AWS applications for your users

119. During regional AWS failures, your API requires the ability to remain online. Stateless API is stored, and you have to add information from other sources as you do not have a DB. How can this Uptime goal be achieved simply but efficiently?

A. Create a Route53 Latency Based Routing Record with Failover and point it to two identical deployments of your stateless API in two different regions. Make sure both regions use Auto-scaling Groups behind ELBs

B. Create a Route53 Weighted Round Robin record, and if one region goes down, have that region redirect to the other region
C. Use an ELB and a cross-zone ELB deployment to create redundancy across datacenters. Even if a region fails, the other AZ will stay online
D. Use a CloudFront distribution to serve up your API. Even if the region your API is in goes down, the edge locations CloudFront uses will be fine

Answer: A
Explanation: Failover routing lets you route traffic to a resource when the resource is healthy or to a different resource when the first resource is unhealthy. The primary and secondary resource record sets can route traffic to anything from an Amazon S3 bucket that is configured as a website to a complex tree of records.

120. A serverless architecture is included the AWS API Gateway, AWS Lambda, and AWS DynamoDB and experienced a large increase in traffic to a sustained 3000 requests per second, which increased the failure rate. The request on the operation lasts for 500 milliseconds on average. The DynamoDB table did not exceed 50% throughput provision, and primary keys are assigned correctly. What can be the reason for failure?

A. Your API gateway deployment is throttling your request

B. Your AWS API Gateway deployment is bottlenecking on request (de)serialization
C. You did not request a limited increase on concurrent Lambda function executions
D. You used consistent read requests on DynamoDB and are experiencing a semaphore lock

Answer: C

Explanation: Every Lambda function is associated with a fixed amount of allocated resources regardless of memory allocation, and each of the functions is allocated with a fixed amount of code storage per function.

121. A CI is needed to build the AMIs with the pre-installed images on every new code push, and you need to do this at a lower cost. How can you do this?

A. Have the CI launch a new on-demand EC2 instance when new commits come in, perform all instance configuration and setup, then create an AMI based on the on-demand instance

B. Bid on spot instances just above the asking price as soon as new commits come in, perform all instances configuration and setup, then create an AMI based on the spot instance
C. Purchase a Light Utilization Reserved Instance to save money on the continuous integration machine. Use these credits whenever you create AMIs on instances
D. When the CI instance receives commits, attach a new EBS volume to the CI machine. Perform all setup on this EBS

volume so that you do not need a new EC2 instance to create the AMI

Answer: B

Explanation: You can bid on spare Amazon EC2 computing capacity in Amazon EC2 Spot instances. Because spot instances are often available at a discount on demand, you can reduce app running costs substantially, increase the computer capacity and performance of your application for the same budget, and allow new types of cloud computing applications.

122. The operation and development team wants a place where it can show both the operation system and application logs. How can you activate this service using AWS? (Choose 2)

A. Using AWS CloudWatch and configuration management, set up remote logging to send events through UDP packets to CloudTrail

B. Using configuration management, set up remote logging to send events to Amazon Kinesis and insert these into Amazon cloud search or Amazon RedShift, depending on the available analytic tool

C. Using AWS CloudFormation, create a CloudWatch Log, log group and send the operating system and application logs of interest using the CloudWatch logs agent

D. Using AWS CloudFormation, merge the application logs with the operating system logs, and use IAM roles to allow both teams to have access to view console output from Amazon EC2

Answer: B and C

Explanation: Amazon CloudWatch logs are used to monitor, store or access your log files from Amazon EC2 instances, CloudTrail, and other sources. You can also retrieve the associated log data from the CloudWatch Log.

123. The development team is using an access key to develop an application that can access S3 or DynamoDB. A new security policy declared that the credentials should not be older than two months and must be rotated. How can this be achieved?

A. Use a script, which will query the data keys that are created. If older than two months, delete them and create new keys

B. Use the application to rotate the keys in every two months via the SDK
C. Delete the user associated with the keys after every two months. Then recreate the user again
D. Delete the IAM role associated with the keys after every two months. Then recreate the IAM role again

Answer: A

Explanation: To get the control keys, you can use CLI command-list-access keys. The CreateDate of the key can be restored as well. When CreateDate is older than two months, it is deleted. Using the command CLI return-list-access key, data about the Access ID for the particular IAM user is retrieved. The operation returns with the empty list if there is none.

124. There is an application deployed using Elastic Beanstalk. Sam has to deploy a new application and ensure that the Elastic Beanstalk has been detached from the current instance and then re-attached to the new instance. But the new instances are not receiving any kind of traffic. What is the case?

A. The instances are of the wrong AMI hence, they are not being detected by the ELB

B. You need to create a new Elastic Beanstalk application because you cannot detach and then reattach instances to an ELB within an Elastic Beanstalk application
C. The instances needed to be reattached before the new application version is deployed
D. It takes time for the ELB to register the instances hence, there is a small time frame before your instances can start receiving traffic

Answer: D

Explanation: Before the traffic starts receiving on an EC2 instance, instances are checked by the ELB health checks, and if the health checks are

successful, the EC2 instance changes their state to an In-service state, and then the instances start receiving the traffic.

125. James configures AWS Inspector, to continuously evaluate EC2 instances (both Linux and Windows) to determine if there are security-related faults and then fix potential issues in time to improve the security of applications deployed on the company's AWS platform. When he defined a new AWS Inspector console assessment target, he chose all AWS EC2 instances. The option to install AWS inspector agents in all instances has also been selected. To produce the inspector's evaluation report, what conditions must be met? (Choose 2)

A. The EC2 instances should have a role to allow SSM Run Command

B. The EC2 instances need to configure an IAM role to have the AWS Inspector full access
C. All EC2 instances need to have the AWS Systems Manager (SSM) Agent installed
D. The security group in the EC2 instances should allow SSH port 22
E. All EC2 instances should have AWS CLI commands pre-installed

Answer: A and C

Explanation: Amazon Inspector is an AWS tool that uses AWS-managed rule packages to perform a security analysis of Amazon EC2 instances. The AWS Inspector agents must be installed first in order for AWS Inspectors to work correctly. You can also select the "Install Agents" option so that the agents are automatically installed. In order to install Inspector agent, SSM Run command is used. For that, IAM role is needed at EC2 instance and all instances must have AWS CLI command pre-installed.

126. A large number of EC2 instances in various AWS accounts have been maintained by a company. Some instances have been started and no longer used for testing purposes. To save cost, they need to work out an approach to quickly identify the EC2 instances that have a low

utilization rate such as daily CPU utilization is 10% or less for several days. What is the best method to choose from the following?

A. In CloudWatch Logs, configure a filter to check the usage rate of EC2 instances. Create an alarm if the utilization is low

B. In Trusted Advisor, check the status of Low Utilization Amazon EC2 Instances, which is part of Cost Optimization Checks
C. Create a Lambda function that checks the CPU utilization for each instance and triggers an SNS notification if the average CPU utilization rate is low
D. In CloudWatch Metrics, for each instance, create an alarm and trigger an SNS notification when CPU utilization is below 10%

Answer: B

Explanation: As they want a method that quickly identifies instances with low utilization, repeatedly checking each instance should be avoided. By using AWS Trusted Advisor, you can easily identify the instance who has low utilization under Cost Optimization Checks.

127. There is a system, which automatically provisions EIPs to EC2 instances on boot in VPC. The system provisions the whole VPC and stack at once, and you have two of them per VPC. You attempt to create a development environment that failed on your new AWS account after successfully creating a staging and production environment in the same region. What is the cause behind this?

A. You did not choose the development version of the AMI you are using

B. You did not set the development flag to true when deploying EC2 instances
C. You hit the soft limit of 2 VPCs per region and requested a third
D. You hit the soft limit of 5 EIPs per region and requested a 6th

Answer: D

Explanation: By default, AWS accounts are limited to 5 Elastic IP addresses per region. You can hit a maximum of 5 EIPs per region as internet addresses (Public (IPv4)) are a scarce public resource. We strongly encourage you to use Elastic IP address primarily in case of instance failure to revert the address to another instance and to use DNS hostnames for all other communication internodes.

128. How can Chris pass queue messages that are 1 GB each?

A. By using AWS EFS as a shared pool storage medium and storing filesystem pointers to the files on disk in the SQS message bodies

B. By using SQS's support for message partitioning and multi-part uploads on Amazon S3
C. By using the Amazon SQS Extended Client Library for Java and Amazon S3 as a storage mechanism for message bodies
D. By using Kinesis as a buffer stream for message bodies. Store the checkpoint ID for the placement in the Kinesis Stream in SQS

Answer: C

Explanation: With Amazon S3, you are able to manage Amazon SQS messages. This is particularly useful for storing and consuming messages up to 2 GB of message size. Use the Amazon SQS Extended Server Library for Java to handle Amazon SQS messages using Amazon S3. You can use this library in particular to:

- Delete the message object from S3
- Get the message object from S3
- Whether on Amazon S3 messages are stored or only if the size of a message exceeds 256 KB
- Send a message, which refers to an object that has been saved in an Amazon S3 bucket

129. Nick and his team examine existing AWS tools to help them better understand where and how AWS saves and accesses sensitive

information. AWS Macie can satisfy the need to analysis, classify and protect data using machine learning. It also offers a dashboard to view various key points of interest such as high-risk S3 items and complete user sessions. Which AWS products are AWS Macie data sources? (Choose 2)

A. AWS S3 Bucket

B. AWS Config
C. AWS CloudWatch
D. AWS CloudTrail
E. AWS EBS Volume

Answer: A and D

Explanation: In order to protect data stored in S3, you can use Amazon Macie, which uses ML. For Amazon Macie, AWS CloudTrail is the data source because it contains API calls, which are provided to Macie for analyzing. AWS S3 bucket is also a data source because when you configure Macie with AWS S3 bucket, then its objects will be classified and monitored.

130. A company needs to build a layer in software stack on AWS that needs to be able to scale depending on demand as quickly as possible. The code is running on an EC2 instance in the Auto-scaling group with ELB. Through which deployment method can this be done?

A. Create a new Auto-scaling Launch Configuration with UserData scripts configured to pull the latest code at all times

B. Create a Dockerfile when preparing to deploy a new version to production and publish it to S3. Use UserData in the Auto-scaling Launch configuration to pull down the Dockerfile from S3 and run it when new instances launch
C. Bake an AMI when deploying new versions of code, and use that AMI for the Auto-scaling Launch Configuration
D. SSH into new instances that come online, and deploy new code onto the system by pulling it from an S3 bucket, which is populated by code that you refresh from source control on new pushes

Answer: C

Explanation: As they want to provide instances as quickly as possible, then it is better to choose the creation of AMI rather than defining it in user data. In AMI, you define the information that is needed for the launching of instances. With AMI, you can launch as many instances as you need.

131. The app is installed in EC2 and sends AWS CloudWatch with customized metrics. Based on these metric report data, you have configured many AWS CloudWatch alarms. You find that sometimes there are numerous warnings of CloudWatch with the status of "INSUFFICIENT DATA". With your team of developers, you have agreed that some metric data are only produced intermittently by design. What should you do to ignore these warnings of "INSUFFICIENT DATA"?

A. Create a CloudWatch Event. When an “INSUFFICIENT_DATA” CloudWatch alarm appears, use an SNS to notify the team to react accordingly

B. Create a Lambda function, which calls CloudWatch Alarm API to check the reason for “INSUFFICIENT_DATA” and modify the alarm state to “OK” if there is no data received from EC2 instance during that time
C. Configure the CloudWatch alarms to change the state of “INSUFFICIENT_DATA” to “OK” after 5 minutes
D. Configure these CloudWatch alarms to treat missing data points as “ignore” so that “INSUFFICIENT_DATA” does not show up

Answer: D

Explanation: If you configure the missing data as ignore, then that data does not cause INSUFFICIENT_DATA alarm.

132. For a new application, Daniel has created a new Auto-scaling group that sets the minimum capacity to 1 and the maximum capacity to 20. ASG has been running smoothly for two weeks, though some instances have not been successfully resolved recently. If an instance fails to terminate, Daniel must work out a way that notifies his team

via email. What is the easiest way to do that?

A. Create a CloudWatch alarm based on the metric of “Terminate Failed”. Send a notification to an SNS topic when the alarm state is “ALARM”

B. Configure the Auto-scaling group to send a notification to an SNS topic whenever instances fail to terminate
C. Configure the related Auto-scaling configuration to send a notification to an SNS topic whenever instances fail to terminate
D. Create a CloudWatch Event. When an event of “Terminate Unsuccessful” happens, invoke a Lambda function to notify the team

Answer: B

Explanation: With ASG, you can configure AWS SNS to send a notification whenever scaling is performed. You can configure SNS topic whenever instances fail to terminate via ASG console.

133. There is a Java program, which deploys an Auto-scaling group in AWS with EC2 instances that are m5.large type instances. The application's response time has recently been increased due to several new functions because of the high rate of use of the CPU. To resolve this issue, you must change the EC2 instance form to c5.large. What is the best way to do so?

A. Select the Auto-scaling launch configuration and choose “Actions -> Change Instance Type”. Modify the instance type to c5.large accordingly

B. In AWS EC2 console, stop the EC2 instances, modify the instance type to c5.large in “Instance Settings -> Change Instance Type” and then start the instances
C. Select the relevant launch configuration and choose “Actions -> Copy launch configuration”. Modify the instance type accordingly in the new configuration. Select the new launch configuration for the Auto-scaling group

D. In AWS console, edit the Auto-scaling group by modifying the instance type from m5.large to c5.large

Answer: C

Explanation: One main feature of Auto-scaling is that it is not modifiable after creation. The best practice is to use an existing configuration as a base for a new setup to change the launch configuration and to update ASG in order to use the new configuration. As the user can pick the current settings, copy the settings to new settings and change the instance type. Users do not have to set everything up from scratch with this approach.

134. For building a high score game table in DynamoDB that will store each user's highest score of players, what will be the DynamoDB structure?

A. Game ID as the hash key, highest score as the range key

B. Highest score as the hash/only key
C. For each game within many games. And each of which has similar usage and the same number Game ID as the hash/only key
D. Game ID as the range/the only key

Answer: A

Explanation: It is best to choose the hash key as the column, which has a wide range of values. You need to sort with the highest score so make the highest score as a sort key.

135. You have 10% of the written and 90% of readings in your web application. All requests are currently being served to an AWS ELB, which is in the front route of the EC2 Auto-scaling group via a Route53 Alias Record. When traffic spikes occur during certain news events, many more people ask to read the same data from your application, all at the same time, making your system become extremely costly. What can you do to reduce costs and scale spikes in the simplest and cheapest way?

A. Create an S3 bucket and asynchronously replicate common requests responses into S3 objects. When a request comes in

for a pre-computed response, redirect it to AWS S3

B. Create another ELB and Auto-scaling group layer mounted on top of another system, adding a tier to the system. Serve most read requests out of the top layer
C. Create a Memcached cluster in AWS ElastiCache. Create cache logic to serve requests, which can be served late from the in-memory cache for increased performance
D. Create a CloudFront distribution and direct Route53 to the distribution. Use the ELB as an origin and specify cache behaviors to proxy cache requests, which can be served late

Answer: D

Explanation: To provide strong reads for your application, use the Cloudfront distribution. Check how long your objects stay in a CloudFront cache before CloudFront transmits another request to your origin. You can create a zone apex record to point to the Cloudfront distribution. The reduction of duration enables dynamic content to be served. Increasing the time means improving the performance of your users because your objects will be served from the edge cache. A longer life also lowers the load on your source.

136. A small IT company has an operating budget for AWS infrastructure that is minimal and therefore spot EC2 instances are always recommended. A new application is introduced through a load balancer and Auto-scaling group. Since this software controls the authentication of all other goods of the business, only spot instances are unacceptable in this regard. According to you, what is the best way to make the Auto-scaling group possible with a combination of on-demand and spot instances?

A. It is impossible to create a combination of On-Demand and Spot instances for this case. Only one type is allowed for an ASG configuration

B. In Auto-scaling Launch Configuration, configure a suitable On-Demand/Spot percentage and then create the ASG with this Launch Configuration

C. Create the ASG by a Launch Template and configure the On-Demand/Spot percentage accordingly
D. Create two ASGs. One for On-Demand instances and one for Spot instances

Answer: C

Explanation: You have two methods of configuring an Auto-scaling group: Launch Template and Launch Configuration.

Launch Configuration does not support the combination of On-Demand and Spot instances. While Launch Template is flexible to support combination instances types and On-Demand and Spot Pricing Options.

137. Your Auto-scaling group scales too quickly, too much, and scales when the traffic is decreasing. What should you do to fix this?

A. Set a longer cooldown period on the group, so the system stops overshooting the target capacity. The issue is that the scaling system does not allow enough time for new instances to begin servicing requests before measuring aggregate load again

B. Calculate the bottleneck or constraint on the computer layer, then select that as the new metric, and set the metric thresholds to the bounding values that begin to affect response latency
C. Raise the CloudWatch alarm thresholds associated with your Auto-scaling group, so the scaling takes more of an increase in demand before beginning
D. Use larger instances instead of lots of smaller ones, so the group stops scaling out and wasting resources as the OS level since the OS uses a higher proportion of resources on smaller instances

Answer: B

Explanation: In the ideal case, the right metric is not used to scale up and down. So in order to fix the issue, create a custom metric of that bottleneck.

138. George has created a DynamoDB table called "Global Temperature" to monitor the highest/lowest temperatures in many cities of different countries. The things in the table had a partition key

(CountryId) listed and there was no sort key. Recently, more features have been created and some queries need to be conducted based on a new partition key (CityId). How can George do this?

A. By adding a Global Secondary Index with a partition key as CityId and a sort key as HighestTemperature

B. By adding a Global Secondary Index with partition key as CityId and another Global Secondary Index with partition key as HighestTemperature. Because the primary index only uses a simple primary key (partition key), the Secondary Index can only have one partition key as well
C. By adding a Local Secondary Index with a partition key as CityId and a sort key as HighestTemperature
D. By modifying the existing primary index with partition key as CityId and sort key as HighestTemperature

Answer: A

Explanation: It is normal for applications to have various queries by using different attributes as Query criteria. More global secondary indexes are necessary in this case. So Secondary Index is the best option to choose as it is comprising of Partition and Sort key.

139. On a new feature, Tony is currently working on to build with an RDS MySQL server. The feature uses the database to store customer data for email transactions. Automated snapshots are designed and function correctly for the server. A new snapshot has to be exchanged and used in another AWS account due to business requirements. How can Tony share this RDS snapshot automatically?

A. There is no way to share snapshots to another account for automated ones

B. Create a manual snapshot for the RDS database and make sure that it is encrypted. Then share the encrypted snapshot to another account via AWS console or CLI
C. Select the automated snapshot in the AWS console. Share the

snapshot by “Actions -> Share Snapshot”

D. Firstly, make a copy of the automated snapshot to turn it into a manual version. Then share the copy with the other AWS accounts

Answer: D

Explanation: You cannot share the automated snapshot with other AWS account but after the snapshot is copied to the manual snapshot, you can use the Share snapshot option from the Actions to share it on another account.

140. An organization has multiple applications in the AWS account, and wants to identify the cost per month to operate for a good understanding of the business as it does not want to expend initial development time. What can be done to achieve this?

A. Use the AWS Price API and constantly running resource inventory scripts to calculate the total price based on the multiplication of consumed resources over time

B. Use AWS Cost Allocation Tagging for all resources, which support it. Use the Cost Explorer to analyze costs throughout the month

C. Use custom CloudWatch Metrics in your system, and put a metric data point whenever the cost is incurred

D. Create an automation script, which periodically creates AWS Support tickets requesting detailed intra-month information about your bill

Answer: B

Explanation: By using the tag on resources you can make resources more organized. To track your AWS costs in a detailed way, you can use tags to organize your resources and costs allocations. After you activate the cost allocation tags, you will be able to categorize AWS and track your AWS costs by using cost allocation tags in order to organize your resource costs in your cost allocation report. AWS provides two types of cost allocation tags, an AWS generated and user defined. AWS defines, creates, and applies the AWS-generated tag for you, and you define, create, and apply user-defined tags. You must activate both types of tags separately before they can appear

in Cost Explorer or on a cost allocation report.

141. During total regional AWS failure, Christopher needs API backup by DynamoDB to stay online. Christopher can only tolerate a few minutes of failure or slowness, but the system should recover those minutes with normal operation. How can he achieve this?

A. Set-up DynamoDB cross-region replication in a master standby configuration, with a single standby in another region. Create an Auto-scaling group behind an ELB in each of the two regions for your application layer, in which DynamoDB is running in. Add a Route53 latency DNS record with DNS failover, using the ELBs in the two regions as the resource records

B. Set-up a DynamoDB multi-region table. Create a cross region ELB pointing to a cross-region Auto-scaling group, and direct a Route53 latency DNS record with DNS failover to the cross-region ELB
C. Set-up DynamoDB cross-region replication in a master standby configuration, with a single standby in another region. Create a cross-region ELB pointing to a cross-region Auto-scaling group, and direct a Route53 latency DNS record with DNS failover to the cross-region ELB
D. Set-up a DynamoDB global table. Create an Auto-scaling group behind an ELB in each of the two regions for your application layer, in which the DynamoDB is running in. Add a Route53 latency DNS record with DNS failover, using the ELBs in the two regions as the resource records

Answer: D

Explanation: DynamoDB global tables provide fully-managed solutions for the deployment of multi-region, multi-master database. And you do not need to maintain and build your own replication solution in DynamoDB global table. You define the region in AWS where the table is to be accessible when creating a global table. DynamoDB carries out all the tasks required in order to construct identical tables and distribute ongoing data changes in these regions.

142. To get CloudFormation stack status updates to show up in a continues delivery system as close to real time as possible, what should you do?

A. Subscribe your continuous delivery system to an SNS topic into which you also tell your CloudFormation stack to publish events

B. Use a long-poll on the resources object in your CloudFormation stack and display those state changes in the UI for the system
C. Use a long-poll on the List stacks API call for your CloudFormation stack and display those state changes in the UI for the system
D. Subscribe your continuous delivery system to an SQS queue that you also tell your CloudFormation stack to publish events into

Answer: A

Explanation: Through monitoring the events of the stack, you will follow the progress of a stack update. The Events table displays every major step in creating and updating the stack, sorted with the latest events, by the time of each event. An UPDATE IN PROGRESS event will mark the beginning of the stack update continuing. When you call CreateStack, use NotificationARNs.member to push stack events into SNS in real-time.

143. George builds an application of photo posting and then images of this application are stored in S3. Now, he wants a system that is simple and cost effective for the application. From the following option, which will be helpful in implementing authentication and authorization to build photo sharing application?

A. Use JWT or SAML compliant systems to build authorization policies. Users log in with a user name and password and are given a token they can use indefinitely to make calls against the photo infrastructure

B. Use the AWS API gateway with a constantly rotating API key

to allow access from the client-side. Construct a custom build of the SDK and include S3 access in it

C. Build the application out using AWS Congito and web identity federation to allow users to log in using Facebook or Google accounts. Once they are logged in, the secret token passed to that user is used to directly access resources on AWS, like AWS S3
D. Create an AWS Auth service domain and grant public sign up and access to the domain. During set-up add at least one major social media site as a trusted identity provider for users

Answer: C

Explanation: You can easily add user sign-up and sign-in, and track your mobile and web app permissions through Amazon Cognito. Inside Amazon Cognito, you can create your own user directory. Use SAML to identify solutions or use your own identity system to authenticate users via social identity providers like Facebook, Twitter or Amazon. Amazon Cognito also enables you to locally store data on devices of users, enabling your apps to operate even when devices go offline. You can then sync data across devices so that you have consistent app experience, regardless of your phone.

144. Your team will start continuous delivery using CloudFormation so that entire, versioned stacks or stack layers can be created and deployed automatically. You have a mission-critical system of 3 levels. What is NOT the best practice for the use of CloudFormation in continuous delivery?

A. Use the AWS CloudFormation validate template call before publishing changes to AWS

B. Use CloudFormation to create a brand new infrastructure for all stateless resources on each push, and run integration tests on that set of infrastructure
C. Parametrize the template and use mappings to ensure your template works in multiple regions
D. Model your stack in one template, so you can leverage CloudFormation’s state management and dependency resolution to propagate all changes

Answer: D

Explanation: CloudFormation's best practices are "created a nested stack" and "re-use templates".

When your infrastructure is growing, common patterns emerge in each template, where you declare the same components. You can separately create templates for these common components. Thus, you can mix and match the various templates, but you can create a single unified stack using nested stacks. Stacks that are nested are stacks that create additional stacks. Use AWS::CloudFormation::Stackresource to link other templates in your template to build nested stacks.

You can reuse your templates to replicate your infrastructure in several environments after you have established your stacks and resources.

145. To deploy multiple stacks of AWS in a repeatable manner in multiple environments, you selected CloudFormation. Now you found that there is a type of resource that you need to create and model, but it is unsupported by CloudFormation. What is the strategy to overcome this challenge?

A. Use a CloudFormation custom resource template by selecting an API call to proxy for create, update and delete actions. CloudFormation will use the AWS SDK, CLI, or API method of your choosing as the state transition function for the resource type you are modeling

B. Create a CloudFormation custom resource type by implementing create, update and delete functionality, either by subscribing a custom resource provider to an SNS topic or by implementing the logic in AWS Lambda

C. Submit a ticket to the AWS Forums. AWS extends CloudFormation resource types by releasing tooling to the AWS labs organization on GitHub. Their response time is usually one day, and they complete requests within a week or two

D. Instead of depending on CloudFormation, use Chef, Puppet or Ansible to author heat templates, which are declarative stack resource definitions that operate over the open stack hypervisor

and cloud environment

Answer: B

Explanation: Custom resources enable you to write custom provision logic in the templates that AWS CloudFormation can run anytime you create, update or delete stacks. When you have changed your custom resource for example, you may wish to include resources that are not available as types of AWS CloudFormation, these resources can be included by using a customized resource. In that way, all your related resources can be managed in one stack.

AWS::CloudFormation::CustomResource or Custom::String resource type to define custom resources in your templates. Custom resources require one property: the service token, which specifies where AWS CloudFormation sends requests to, such as an Amazon SNS topic.

146. Christian met his operation team to discuss last month's data. During the meeting, he realized that three weeks ago, his monitoring system, which pings over HTTP from outside the AWS recorded a large spike in latency on his 3 tier web service API. DynamoDB is used for database layer, EBS, ELB, EC2 for the business logic tiers and SQS, EC2, and ELB for the presentation layer. Which technique will not figure out what happened?

A. Review CloudWatch metrics for one-minute interval graphs to determine which component(s) slowed the system down

B. Check your CloudTrail log history around the spikes time for any API calls that caused slowness
C. Review your ELB access logs in S3 to see if any ELBs in your system saw the latency
D. Analyze your logs to detect bursts in traffic at that time

Answer: A

Explanation: CloudWatch retention is:

- Data point with less than 60 seconds are available for 3 hours
- 60 seconds are available for 15 days
- 300 seconds are available for 63 days

- 3600 seconds are available for 455 days

Data points that are initially published with a shorter period are aggregated together for long-term storage. For example, if you collect data using a period of 1 minute, the data remains available for 15 days with a 1-minute resolution. After 15 days, this data is still available, but is aggregated and is retrievable only with a resolution of 5 minutes. After 63 days, the data is further aggregated and is available with a resolution of 1 hour

147. A vendor requires access to the S3 bucket in your account. The vendor already has an AWS account. How can you provide access to the bucket?

A. By creating a new IAM user and granting the relevant access to the vendor on that bucket

B. By creating a new IAM group and granting the relevant access to the vendor on that bucket
C. By creating an S3 bucket policy that allows the vendor to read from the bucket from their AWS account
D. By creating a cross-account role for the vendor account and grant that role access to the S3 bucket

Answer: D

Explanation: You can share resources in one account with the users in a different account. By cross-account access, you do not need to create individual IAM users in each account. The users do not have to sign out and then sign in in another account to get access to the AWS resources. You can use IAM roles and STS to set-up cross account access.

148. An application is deployed, which uses Auto-scaling for launching the new instances. To change the instance type of the new instances, which of the action is deployed?

A. Using Elastic Beanstalk to deploy the new application with the new instance type

B. Creating a new launch configuration with the new instance type
C. Using CloudFormation to deploy the new application with the

new instance type

D. Creating new EC2 instances with the new instance type and attach it to the Auto-scaling group

Answer: B

Explanation: Create a new configuration, attach it with the existing Auto-scaling group and then terminate the running instances.

149. There is an application hosted on AWS, on EC2 instance behind a load balancer. You add new features on this application, which causes the sites to slow down, and now you are receiving complains. How can you recover from this issue?

A. By using CloudTrail to log all the API calls, and then traversing the log files to locate the issue

B. By using CloudWatch and monitoring the CPU utilization to see the times when the CPU peaked

C. By creating some custom CloudWatch metrics, which are pertinent to the key features of your application

D. By reviewing the Elastic Load Balancer logs

Answer: C

Explanation: The issue could be relevant to the few features. Enabling CloudWatch to monitor all the API calls of all services will not benefit the cause. The monitoring of CPU utilization will re-verify that there is some issue but will not resolve the issue. ELB logs do the same things. So, you need to create a custom metric in CloudWatch for this requirement.

150. A team is assisting you in creating a new application with the AWS Aurora data base. The application's users are mainly from Europe and North America. You suggest configuring a global Aurora database. The primary server is in the eu-west-1 area and the secondary one in the us-east-1 zone. What are the benefits of Amazon Aurora Online Database? (Choose 2)

A. There is no charge for the replicated write I/Os between the primary region and each secondary region

B. For disaster recovery, the secondary cluster can be easily promoted to allow full read and write operations
C. The clusters in both primary and secondary regions have the same read & write configured capacities
D. The cluster in the secondary region us-east-1 enables low-latency reads
E. Aurora Global Database is available for either Aurora MySQL or Aurora PostgreSQL

Answer: B and D

Explanation: In order to create a Global DB, you need to select the location Global. Since users provided read-only services from the secondary AWS Region server, there is no write for the secondary region. With this Global DB when required, the database in the secondary AWS Region can be promoted to take full workloads within a minute.

151. Louis has used AWS CodeDeploy to deploy the latest software release to several AWS EC2 instances. In this CodeDeploy application, there are two Deployment groups called Stage and Prod. At the moment, the only difference between Stage and Prod is the logging level, which can be configured in a file. What is the most efficient way to implement the different logging levels for the two Deployment groups?

A. Create two script files. One version has its logging level setting for “Stage”. And the other one has its logging level setting for “Prod”. Modify the Deployment Group configurations to use the correct script file

B. In the hook script file, use the environment variable DEPLOYMENT_ID to determine, which Deployment Group it is. Then modify the logging level accordingly in the script
C. For the script file in BeforeInstall hook, use the environment variable DEPLOYMENT_GROUP_NAME to determine the Deployment Group. Then modify the logging level accordingly in the script
D. Create two source versions of script files in BeforeInstall hook. One version has its logging level setting for Deployment Group

"Stage". And the other one has its logging level setting for "Prod". Choose the relevant source files when creating a new deployment

Answer: C

Explanation: Since the only difference between deployment groups is the logging level, the environment variable to decide the deployment group should use the same number of source files. Because the environment parameter to evaluate the Deployment Group is DEPLOYMENT GROUP NAME, in the script file, the following code can be used:

["$DEPLOYMENT_GROUP_NAME" == "Staging"]

After this, modify the logging level for Staging Deployment Group.

152. The company you work for recently began using the AWS system and used a basic support package for its AWS account. As AWS resources are limited, your director requested to develop a solution to alert him to the low utilization rate of EC2 assets. You intend to use AWS CloudWatch Events to define the lambda feature as the target for these tools. The Lambda feature can provide the director with personalized text. How are you to combine these two choices to fulfill your requirements? (Choose 2)

A. Create a new CloudWatch Events rule with the event source as a Trusted Advisor. The event type is "Check Item Refresh Status". Select the Lambda function as the target

B. Create a CloudWatch rule with the event source as a Trusted Advisor. The event type is "All Events". Select the Lambda function as the target
C. Upgrade the AWS account support plan to "Business" to access all Trusted Advisor checks
D. Upgrade the AWS account support plan to "Developer" to access all Trusted Advisor features
E. Upgrade the AWS account support plan to "Enterprise" to activate all cost related features

Answer: A and C

Explanation: The Business and Enterprise Support Plan should be used for

the Cost Optimization Checks for trusted advisors. The Trusted Advisor Cost Optimization Checks can capture low-cost EC2 instances. To test the trusted advisor activities, the client just has to create a CloudWatch event policy and provide notification with Lambda.

153. In order to activate CloudTrail, the company wants to increase visibility in AWS user and resource activities. The trail was set up only to collect activities in the ap-south-1 region since most resources were used. However, customers are concerned that the data in the CloudTrail could be lost if a disaster occurs in the ap-south region. What measures should you take to tackle the problem?

A. Create a Lambda Function, which can read files in the trail S3 bucket and copy over the log files to an S3 bucket in another region

B. Enable Cross-Region Replication for the trail S3 bucket, which automatically copies objects in different AWS Region
C. In CloudTrail console, add another S3 bucket in a different region as the target for the trail
D. Enable the encryption with SSE-KMS for the S3 bucket of the CloudTrail

Answer: B

Explanation: When failure occurs in the AWS region, you need to consider that there is no data loss in CloudTrail. So with Cross-Region Replication, your data asynchronously copies to another region.

154. You have to set up a solution that incorporates single sign-on from your corporate AD or LDAP directory and does not allow access to each user to a designated user folder in a bucket. Choose 3 answers to fulfill the given scenario.

A. Set up a federation proxy or identity provider
B. Tag each folder in the bucket
C. Configure IAM role
D. Use AWS STS service to generate temporary tokens
E. Set up a matching IAM user for every user in your corporate

directory that needs access to a folder in the bucket

Answer: A, C, and D

Explanation: Firstly, the access request is sent to the identity provider, which directs the request to the corporate identity store, which authenticates the user and sends the request to STS, which issues the temporary token to the user, and then the user can login to the console and get access of the desired bucket.

155. Specific attention is needed to ensure that the security vulnerability is not present for its AWS services. A rule is provided in AWS Config to test if the assets of AWS are always according to expectations. The policy is very complex and the current AWS-managed rule is not applicable to the needs of the organization. What actions can this requirement achieve in combination? (Choose any 2)

A. In AWS Config, add a custom rule that runs every hour and sends a message to an SQS queue

B. Create a t2.micro EC2 instance to implement the custom policies to check if AWS resources are not exposed to security issues. The instance also listens to an SQS queue and starts processing whenever there is a new message in the queue
C. In AWS Config, add a custom rule and specify the ARN of an AWS Lambda function that checks AWS resources
D. Create an SNS topic and subscribe to the topic with an email notification to the team member in case there is a security issue in AWS resources
E. Create an AWS Lambda function, which contains the logic of the custom rule to evaluate whether the AWS resources are compliant

Answer: C and E

Explanation: AWS config has the rules or custom rules managed by AWS. A Lambda function can be set with the Config-Rule-Triggered blueprint. A Lambda ARN can be configured when a custom rule is created.

156. You have an instance in Auto-scaling group, in which lifecycle

hooks are enabled. Due to lifecycle hooks, initially, the instance is put into Pending: Wait state, which means that the instance cannot handle the traffic in this state. In wait sate, other scaling actions are also suspended. Later, the instance is put into Pending: Proceed and then it is changed to InService, which means that instances in Auto-scaling group can now serve the traffic, but you have observed that the bootstrapping process finishes earlier than the status Pending: Proceed is updated.

What can you do to check that the status of the instance is updated correctly after the bootstrapping process?

A. Use the complete-lifecycle-action call to complete the lifecycle action. Run this command from Command Line Interface
B. Use the complete-lifecycle-action call to complete the lifecycle action. Run this command from another EC2 instance
C. Use the complete-lifecycle-action call to complete the lifecycle action. Run this command from an SQS queue
D. Use the complete-lifecycle-action call to complete the lifecycle action. Run this command from the Simple Notification Service

Answer: A

Explanation: Use the complete-lifecycle-action command to allow the Auto-scaling Group to continue to launch or terminate the instance after completing your custom action before the time period expires. You can use the following command to specify the lifecycle action token:

Aws auto-scaling complete-lifecycle-action –lifecycle-action-result.

157. The AWS API Gateway and Lambda service both have been implemented by a team. The JSON input data can be saved to the internal database and S3. Some data in the JSON body are not used after the service runs for several years, and should no longer be saved. The Lambda, which is at the backend was modified to only support the current JSON values. Before all users move to the new API, the original API in API Gateway remains used for a while. How should the original API be set to still support the backend of the old application by Lambda?

A. In Integration Response of the API, add a mapping template to remove the obsolete data for the backend Lambda to support the old requests

B. Configure a mapping template in Integration Request to remove the obsolete data so that the original API requests are transformed to be supported by the new Lambda
C. In the original API, add a stage for canary deployment to understand how many users are still using the old JSON format before the original API service is completely removed
D. In the original API, use a new Lambda as an authorizer so that only the requests with valid JSON data can proceed to hit the backend

Answer: B

Explanation: The Lambda backend has already been updated and the original API has been preserved for some time. However, Lambda supports only the new requests without any outdated information within the JSON body, meaning the original requests in the Integration Request have to be mapped to a mapping template. Because the request can be converted to the correct format with a mapping template, nothing has changed from the initial viewpoint of users.

158. You want to design a .Net front end and DynamoDB back end application. You know that the application will run with a heavy load. How will you ensure the scalability of application for minimum DynamoDB database load?

A. By launching DynamoDB in Multi-AZ configuration with a global index to balance writes
B. By using SQS to assist and let the application pull messages and then perform the relevant operation in DynamoDB
C. By increasing the write capacity of DynamoDB to meet the peak loads
D. By adding more DynamoDB databases to handle the load

Answer: B

Explanation: SQS is the best option for scalability. DynamoDB is usually

scalable, messages in SQS can help in the management of the above-mentioned situation due to the cost effective solution's condition. Amazon Simple Queue Service (SQS) is a fully managed service for the communication of message queues between distributed microservices and software components at any scale. It is the best practical design for modern applications.

SQS makes decoupling and coordinating the components of a cloud application simple and cost- effective. You can send, store and receive messages from software components at any volume via SQS without losing messages or demanding the availability of other services at any time.

159. To maintain an API endpoint, your team uses the AWS API Gateway and Lambda functionality. Recently, the Lambda function has been modified and some new data is applied to the response body. The Lambda feature has been updated to support the change and a new API has been added. The team wants to keep API version 1 and version 2 accessible to your API consumers simultaneously. In order to make the version 1 API compliant, how should you change it in the API Gateway?

A. Add a mapping template in the Integration Response of version1 API to remove the new data so that the response is transformed into the original format

B. In the Integration Response of version1 API, add a new Header Mapping to remove the new data, which comes from the backend
C. In the Integration Request of version1 API, add a mapping to inform the backend Lambda to remove the new data in the relevant response
D. Add a mapping template in the Method Response of version1 API to remove the new data in the response to keep the same behaviors as before

Answer: A

Explanation: To change the backend, the Integration Response of the old API could add a mapping model. Because mapping template is suitable for setting the backend response, the response for the original API consumers is

exactly the same.

160. You are working in a company where you have to record all the activities occurring in AWS account and provide the access of the loggings of events to the security officer across all regions in a simple and secure way that no one else would be able to access those events other than the security officer. How will you execute your task? Select the best solution from the given options.

A. Use CloudTrail to log all events to an Amazon Glacier vault. Make sure the vault access policy only grants access to the security officer's IP address

B. Use CloudTrail to log all events to separate S3 buckets in each region as CloudTrail cannot write to a bucket in different regions. Use MFA and bucket policies on all the different buckets
C. Use CloudTrail to log all events to one S3 bucket. Make this S3 bucket is only accessible by your security officer with a bucket policy that restricts access to his user only and also add MFA to the policy for a further level of security
D. Use CloudTrail to send all API calls to CloudWatch and send an email to the security officer every time an API call is made. Make sure the emails are encrypted

Answer: C

Explanation: CloudTrail is used to enable security analysis, track changes to your account and provide compliance auditing. You can log and monitor and retain events relating to API calls throughout your AWS infrastructure continuously with CloudTrail. A history of your account AWS API calls, including API calls via the AWS Management console, AWS SDKs, command line tools, and other AWS services, is provided by the CloudTrail. This history helps security analysis, tracking, and troubleshooting of resources.

161. You are tasked to move an application to the world of AWS Elastic Beanstalk. The OS is Amazon Linux and package managers need 'yum' and 'rubygems' to download all packages required for the

program. In .ebextensions, you have created config files in which you have added a software package section. For the package segment in the configuration folder, what is the correct statement?

A. Elastic Beanstalk only supports one package manager per config file. So at least two config files are needed. One for yum and the other for rubygems

B. Within each package manager, the package installation order is not guaranteed
C. For Amazon Linux Operating System, only the package manager yum is supported in the Package section. Elastic Beanstalk does not support rubygems
D. Only the latest supported version is installed. The version number cannot be specified in the configuration file

Answer: B

Explanation: In the Package manager, the order of package installation is not guaranteed. For example, you define:

Packages:

Yum:

Application 1: []

Application 2: []

In that, it may be possible that application 2 installs first.

162. An IT administrator has the responsibility to create a development environment, which would confirm the LAMP development stack. Whenever a new instance is launched, the development team should be updated with the latest version of the LAMP stack. Choose 2 answers, which will meet the requirements in the best way possible.

A. Use the User data section and use a custom script, which will be used to download the necessary LAMP stack packages

B. Create a CloudFormation template and use the cloud-init directives to download and install the LAMP stack packages
C. Create an EBS volume with the LAMP stack and attach it to an

instance whenever it is required

D. Create an AMI with all the artifacts of the LAMP stack and provide an instance to the development team based on AMI

Answer: A and B

Explanation: You can always ensure that the latest version of the LAMP stack is downloaded and given to development teams using user data and cloud-init directive. The AMI's version should always be the same, and you must create an AMI each time you change the version of the LAMP stack.

You can transfer your user data to the instance when you launch an instance in Amazon EC2, which can be used to carry out common automated configuration tasks and even run scripts after the instance is started. The two types of user data that can be passed to EC2 instance are shell scripts and cloud-init directives. You can also transfer that data into a launch wizard either as a file (which can be used to launch instances using the command line tools), as simple text, or as a base64-encoded text (for API calls).

163. If you are asked to build a social media mobile application, which needs permission for every user login and storing their data in DynamoDB. What would you choose to do from the following options in order to grant access to DynamoDB to your application users when required?

A. Create an active directory server and an AD user for each mobile application user. When the user signs in to the AD sign-on, allow the AD server to federate using SAML 2.0 to IAM and assign a role to the AD user, which is the assumed with AssumeRoleWithSAML

B. During the install configuration process, each user creates an IAM credential and assign the IAM user to a group with proper permissions to communicate with DynamoDB.

C. Create an IAM group that only gives access to your application and to the DynamoDB tables. Then, when writing to DynamoDB, simply include the unique device ID to associate the data with that specific user

D. Create an IAM role with the proper permission policy to communicate with the DynamoDB table. Use web identity

federation, which assumes the IAM role using AssumeRoleWithWebIdentity, when the user signs in, grant temporary security credentials using STS

Answer: D

Explanation: To access any AWS service, using a role is the prior way to approach any application whereas web identity federation is used for any web application. To develop a web application, it is recommended that its long term AWS credentials should not be installed or distributed with apps even in an encrypted store that a user downloads. Instead, build your app where the AWS temporary security credentials dynamically web identity federation is required. The temporary credentials mapping to an AWS role only allow executing the tasks required by the mobile app.

164. You are building a pipeline for a new Web application that AWS CodePipeline. The source provider in the source stage is GitHub and its output artifact is designed as WebApp. In the Build process, two simultaneous development activities have been performed and the service providers have used AWS CodeBuild. Which input/output artifact configuration is right for these two build actions?

A. Build Action 1 (input artifact: WebApp, output artifact:WebAppBuild); Build Action 2 (input artifact: WebApp, output artifact:WebAppBuild)

B. Build Action 1 (input artifact: WebApp, output artifact:WebAppBuild1); Build Action 2 (input artifact: WebApp, output artifact:WebAppBuild2)

C. Build Action 1 (input artifact: empty, output artifact:WebAppBuild1); Build Action 2 (input artifact: WebApp2, output artifact:WebAppBuild2)

D. Build Action 1 (input artifact: WebApp, output artifact:WebAppBuild1); Build Action 2 (input artifact: WebApp2, output artifact:WebAppBuild2)

Answer: B

Explanation: To set up input/output artifacts of CodeBuild stage in CodePipeline. One thing to note is that CodeBuild should have at least 1

input artifact. So Option B is correct as it has one input artifact for both stages.

165. A DevOps engineer is constructing a pipeline for an ongoing Java project. The CI/CD software is chosen as AWS CodePipeline. CloudFormation stacks are used during the deployment process. Nonetheless, a variety of conditions must be taken into account. For example, if there is no stack, the stack may be generated. Check whether a stack already exists or should be created and then inspected. The team managers tend to introduce an overviewed state machine and each state executes a Lambda operation. How can this requirement best be fulfilled?

A. Use a shell script in EC2 to interface with AWS MQ service to achieve the function of the state machine. Depending on the running status of AWS MQ service, the script returns the execution result back to CodePipeline

B. Establish several SQS queues to indicate running status. The Lambda function for each state gets the message from the queue and processes the deployment task accordingly. The final result is returned to CodePipeline by a dedicated Lambda
C. Use a Python script in EC2 to record the running status and achieve the state machine feature. It calls different Lambda functions depending on the current status. The script returns the execution status back to CodePipeline
D. Use a Lambda Function to interface with an AWS Step Functions, which implements the workflow-driven state machines for CloudFormation stack deployment. The Lambda Function returns the execution status back to CodePipeline

Answer: D

Explanation: This question requires a visualized state machine. The best tool used to achieve this is AWS Step Functions. It becomes easier for users to understand which tasks were performed and why a state is reached because of the AWS Step features. It can also communicate with Lambda and inform AWS CodePipeline about its status.

166. Your company CTO assigns you that task to connect a MySQL Database to a Wordpress application keeping in ming that the environment must be fault tolerant and highly available. Select 2 options from the following, which would individually play their role to perform the required task.

A. Create a MySQL RDS environment and create a Read Replica

B. Create Multiple EC2 instances in the same AZ. Host MySQL and enable replication via scripts between instances
C. Create a MySQL RDS environment with Multi-AZ feature enabled
D. Create Multiple EC2 instances in separate AZ. Host MySQL and enable replication via scripts between instances

Answer: C and D

Explanation: If you want high availability and fault tolerant environment, the instances must be located in multiple availability zones. Therefore, if you host your own MySQL, ensure that you have instances spread over several AZs.

By using Multi-AZ deployments, Amazon RDS delivers high availability and failover support for DB instances. Amazon's failover technology is used by Multi-AZ deployments for PostgreSQL, MySQL, Oracle, and MariaDB DB instances.

167. A company is designing loosely coupled system. For executing this successfully, which of the following design strategies are ideal? (Choose 2)

A. Having the web and worker roles running on the same set of EC2 instances

B. Having the web and worker roles running on separate EC2 instances
C. Using SQS to establish communication between the web and worker roles
D. Using SNS to establish communication between the web and worker roles

Answer: B and C

Explanation: You can use SQS and separate environments for web and worker processes. Communication between the web and worker roles is managed by the SQS queue.

168. You assist a team in developing a new AWS system CI/CD Pipeline. The pre-requisite is that an orchestration system like origin control, design, and implementation handle the pipeline as a whole. Throughout implementation, various workflow-driven activities are taken into account to ensure that the state machine engine is able to handle these actions correctly. The execution status of the pipeline and state machine should be visualized. What two combined approaches can implement this? (Choose 2)

A. Configure AWS Step Functions State Machine to process the state transition. Return the running state when being asked

B. Create an AWS SWF workflow for the state machine execution. Use workflow decider to control the activity steps
C. Setup AWS CodeStar to manage the whole CI/CD services in AWS
D. Configure an AWS CodeDeploy application to process the continuous deployment tasks
E. Set up an AWS CodePipeline. In its deployment stage, call a Lambda function, which is the trigger for the state machine

Answer: A and E

Explanation: If an orchestration tool is required in AWS, CodePipeline should be considered as the best service to manage source control, building, and deployment. AWS CodePipeline automates the steps towards the continuous publishing of software changes, which is an ideal CI/CD tool. You can also integrate Step Functions with CodePipeline.

With AWS Step Functions, it becomes easier for users to understand which tasks were performed and why a state is reached. It can also communicate with Lambda and inform AWS CodePipeline about its status.

169. An organization has an Auto-scaling group with the following setting: Minimum capacity: 2, Desired capacity: 2 and Maximum

capacity: 4. In the Auto-scaling group, the total number of instances is two that are currently running. You are advised to ensure that no new instances were introduced by the Auto-scaling team for a duration of an hour. Which combination of steps will avoid the release of new instances? (Choose 2)

A. Change the Minimum capacity to 2

B. Suspend the Launch process of the Auto-scaling Group
C. Change the Desired capacity to 4
D. Change the Maximum capacity to 2

Answer: B and D

Explanation: You may pause the creation of new instances temporarily by limiting the maximum capacity to 2 so that the existing instances of 2 will be equal to the maximum limit. Secondly, you can suspend the launch process of the Auto-scaling Group.

170. How will you, as DevOps engineer, automate the creation of EBS snapshot?

A. By using Cloudwatch Events to trigger the snapshots of EBS Volumes

B. By using the AWS CodeDeploy service to create a snapshot of the AWS Volumes
C. By using the AWSConfig service to create a snapshot of the AWS Volumes
D. By creating a powershell script, which uses the AWS CLI to get the volumes and then run the script as a cron job

Answer: A

Explanation: The best thing to do is to use CloudWatch's built- in service as CloudWatch events to automate the creation of EBS snapshots. With Option A, you should only run the power shell script on Windows machines and keep the script itself. And you have the overhead to just run this script with a separate instance.

When you go to CloudWatch events, you can use the Target as EC2 CreateSnapshot API call.

CloudWatch Events provides an almost real-time stream of system events describing changes in Amazon Web Services (AWS) resources. You can match and route events to one or more target functions or streams by using simple rules that you can set up quickly.

171. You are a partner for AWS and assist a customer in the development of applications with AWS Elastic Beanstalk network. In the Amazon Linux system, one significant software will be installed. You also built some config files in the.ebextensions folder because there are several Shell commands to execute before and after the program version is extracted. What parts of the config file can you insert commands to make it work properly? (Choose 2)

A. deploy section

B. container command section
C. file section
D. services section
E. commands section

Answer: B and E

Explanation: The application Elastic Beanstalk can be configured in.ebextensions folder by the configuration of files. Before the application is set up and the application version file is extracted, the commands specified in commands sections are executed. This section is therefore required. The container command main section supports executing commands after extracting an application version file.

172. A DevOps engineer helps the development team build AWS Elastic Beanstalk's Node.js software. The maximum number of Auto-scaling group instances was updated to 10 in AWS console during environment development. In the source bundle, however, the .ebextensions folder contains a configuration file setting the maximum number of instances to 5. Since the limit of 4 is standard for Elastic Beanstalk, how many ASGs can the DevOps engineer build in the AWS console after the environment?

A. 5

B. 6
C. 4
D. 10

Answer: D

Explanation: To modify the configuration in the Elastic Beanstalk, you can use AWS CLI and Console. The priorities are “configurations in AWS console” > “configurations in .ebextensions” > “default value”. Therefore, the maximum value of the ASG will be 10.

173. Allen created a CloudWatch events rule to capture CloudFormation API call:

```
{
  “source”: [
    “aws.cloudformation”
   ],
“detail-type”: [
“AWS API Call via CloudTrail”
],
“detail”: {
   “eventSource”: [
       “cloudformation.amazon.com”
  ]
}
}
```

The Event rule has the Lambda Function as a target for slack channel notifications. Nevertheless, no notification was provided when a CloudFormation stack was modified. Which one of the following could be the cause?

A. Check if CloudTrail logging is turned on in this region. If it is turned off, no AWS API action events can be received

B. The rule should change in the following snippet

```
"source": [
    "aws.cloudtrail"
  ],
"detail": {
   "eventSource": [
       "cloudformation.amazon.com"
  ]
}
}
```

C. CloudTrail can only trace the add and delete for a CloudFormation stack. No API call is recorded for the stack update so no notification was received
D. AWS API Call via CloudTrail cannot be used in a CloudWatch Event rule

Answer: A

Explanation: The event type can be set to 'AWS API Call via CloudTrail' when a CloudWatch Event rule is created. Because the CloudWatch Events Regulations "AWS API Call via CloudTrail" depends on CloudTrail functions, the events rules cannot be triggered by CloudTrail first if CloudTrail is not turned on.

174. A company wants to use a SaaS third-party app running on AWS. The SaaS application must be able to issue several API commands to detect Amazon EC2 resources in the company's account. The company has internal security policies that require external access to its environment and must ensure with the minimum privilege principles and controls, which ensure that the credentials used by the SaaS vendor cannot be accessed by the third party. Which of the following would satisfy all these requirements?

A. Create an IAM user within the enterprise account assigning a user policy to the IAM user that allows only the actions required by the SaaS application. Create new access and secret key for the user and provide these credentials to the SaaS provider

B. From the AWS Management Console, navigate to the Security Credentials page and retrieve the access and secret key for your account
C. Create an IAM role for EC2 instances, assign it a policy that allows only the actions required for the Saas application to work, provide the role ARN to the SaaS provider to use when launching their application instances
D. Create an IAM role for cross-account access allowing the SaaS provider's account to assume the role and assign it a policy that allows only the actions required by the SaaS application

Answer: D

Explanation: Many SaaS platforms provide access to AWS resources via created AWS cross account access. You will see the ability to add a cross-account role if you go to Roles in your identity management.

175. A company wants to access the on-premises LDAP server from an application launched on a VPC. The connection of VPC and on-premises location is established through an IPSec VPN. Choose 2 correct options for application user authentication.

A. The application authenticates against LDAP and retrieves the name of an IAM role associated with the user. The application then calls the IAM Security Token Service to assume that IAM role. The application can use the temporary credentials to access any AWS resources

B. Develop an identity broker that authenticates against LDAP and then calls IAM Security Token Service to get IAM federated user credentials. The application calls the identity broker to get IAM federated user credentials with access to the appropriate AWS service
C. Develop an identity broker that authenticates against IAM Security Token Service to assume an IAM role in order to get temporary AWS security credentials. The application calls the identity broker to get AWS temporary security credentials
D. The application authenticates against LDAP and then calls the IAM Security Service to log in to IAM using the LDAP

credentials. The application can use IAM temporary credentials to access the appropriate AWS service

Answer: A and B

Explanation: If you need an on-premises environment to work with a cloud environment, you usually have two artifacts for authentication:

An Identity Store: This is the on- site store like Active Directory, where information of all the users and groups are stored.

An Identity Broker: It acts as an intermediate agent between the cloud environment and the on-premises location. This facility is provided by a system called Active Directory Federation services in Windows.

The outside user is first authenticated by the Identity broker using active directories. Then the temporary security token is issued to access console or users' access APIs.

176. There are several AWS CloudFormation templates that Harry has to maintain. Two incidents occurred last month where someone modified assets in existing CloudFormation stacks without any alert or notification. This had negative effects on the stacks of possible changes, and the changes were lost when the stacks were reset. This is not compliant with company policy. His team leader has demanded that he alerts the team when drifts occur in the CloudFormation stack. What is the best way to do this?

A. Enable CloudTrail. Use a Lambda function to analyze the CloudTrail logs. Send an email if it is found that the resources in CloudFormation stacks are modified

B. Create a CloudWatch event rule for CloudFormation. If any event happens for CloudFormation, trigger an email notification by SNS
C. Create a rule in AWS Config to evaluate if the stack is considered to have drifted for its resources. If the rule is NON_COMPLIANT, notify the team via an SNS notification
D. Use a Lambda function to check the drift status in each CloudFormation stack every 10 minutes. If there is a drift, send the team an email via AWS SES service

Answer: C

Explanation: There is a rule operated by AWS in AWS Config (cloudformation-stack-drift-detection-check) to see if drift occurs. So, quickly set drift status rule and users can see whether or not the stacks are compliant. Notify everyone then via AWS SES.

177. You are using an EC2 instance, on which web and worker role infrastructure is defined. Jobs sent by the web role are managed through SQS. Now, what way will be suitable to check that the jobs are sent properly by the web role?

A. Use CloudWatch monitoring to check the size of the queue and then scale out SQS to ensure that it can handle the right number of jobs

B. Use CloudWatch monitoring to check the size of the queue and then scale out using Auto-scaling to ensure that it can handle the right number of jobs
C. Use ELB to ensure that the load is evenly distributed to the set of web and worker instances
D. Use Route53 to ensure that the load is evenly distributed to the set of web and worker instances

Answer: B

Explanation: SQS can be used to manage communication between the web and the roles of the worker. The number of messages in the SQS queue can be used to determine the number of instances the Auto-scaling group should have.

178. Jose is working in a company that has to do a major public announcement of a social media site in AWS. The website is running on EC2 instances launched in multiple availability zones with a Multi-AZ RDS MySQL Extra Large DB Instance. The site executes a high number of small reads and writes per second and depends on an eventual consistency model. After complete tests, he observes that there is read contention on RDS MySQL. Which approaches should he choose to meet these requirements? (Choose 2)

A. Deploying ElasticCache in-memory cache running in each availability zone

B. Increasing the RDS MySQL Instance size and implementing provisioned IOPS
C. Adding an RDS MySQL read replica in each availability zone
D. Implementing sharding to distribute the load to multiple RDS MySQL instances

Answer: A and C

Explanation: Enhanced performance and durability for database instances are provided by Amazon RDS read replicas. For read-heavy database workloads, this replication feature makes it easy to scale out beyond a capacity constraint of a single DB instance.

Amazon ElastiCache is used to deploy, run and scale an in-memory data store or cache in the cloud. It provides a high performance web application because it allows you to fetch information from fast, managed, in-memory data stores.

179. A company has given you the task of designing a CloudFormation template. In the template, it must be included that the CloudFormation stack is deleted and a snapshot of the Relational DB is created as a part of the stack. What should you do to complete this task?

A. Create a new CloudFormation template to create a snapshot of the relational database

B. Create a snapshot of the relational database beforehand so that when the CloudFormation stack is deleted, the snapshot of the database is present
C. Use the Update policy of the CloudFormation template to ensure a snapshot is created of the relational database
D. Use the DeletionPolicy of the CloudFormation template to ensure a snapshot is created of the relational database

Answer: D

Explanation: When the resource stack is deleted, you can store or (in some

cases) backup a resource with the DeletionPolicy attribute. For each resource you would like to control, you have to specify a DeletionPolicy attribute. AWS CloudFormation can remove the resource by default if a resource has no DeletionPolicy attribute. Please note that this ability also covers the update operations, which results in the removal of resources.

180. AWS DevOps admins are working on building the infrastructure of the company's development team through CloudFormation, which includes the VPC and networking components, installing a LAMP stack and securing the created resources. For designing this template, choose the best way among the following options.

A. Create multiple CloudFormation templates based on the number of development groups in the environment

B. Create multiple CloudFormation templates for each set of logical resources, one for networking, and the other for LAMP stack creation
C. Create multiple CloudFormation templates based on the number of VPCs in the environment
D. Create a single CloudFormation template to create all the resources since it would be easier from the maintenance perspective

Answer: B

Explanation: Creating multiple CloudFormation templates is one of the examples of nested stacks. When infrastructure grows, common patterns may arise in which each of your templates declare the same components. You can create dedicated templates and separate common components. In this way, you can mix and match various templates, but use nested stacks to create a single stack. To create other stacks with in a stack, one should use nested stacks. AWS::CloudFormation::Stack resource is used in the template to refer other templates to create nested stacks.

181. A company is using OpsWork with several stacks for development, staging, and production to deploy and manage the application. The company wants to start using python instead of Ruby as it already as Ruby on Rail content management platform. Choose

the correct solution to manage the new deployment.

A. Create a new stack that contains a new layer with the Python code. To cut over the new stack the company should consider using Blue/Green deployment

B. Update the existing stack with Python application code and deploy the application using the deploy life-cycle action to implement the application code
C. Create a new stack that contains Python application code and manage separate deployments of the application via the secondary stack using the deploy life-cycle action to implement the application code
D. Create a new stack that contains Python application code and manage separate deployments of the application via the secondary stack

Answer: A

Explanation: Blue/green deployment is the technique for application release by shifting of traffic from different application versions in two identical environments. Blue / green deployments can reduce common risks associated with software deployments, such as down time and rollback.

182. You are responsible for creating a cloud information template to spin resources on your DevOps team's demand. The necessity is to distribute assets in different regions through this cloud creation model. Which of the following aspects will help you develop a model to spin the resources based on region?

A. Use the metadata section in the Cloudformation template, so that based on the relevant region, the relevant resource can be spun up

B. Use the parameters section in the Cloudformation template, so that based on the relevant region, the relevant resource can be spun up
C. Use the mappings section in the Cloudformation template, so that based on the relevant region, the relevant resource can be

spun up

D. Use the outputs section in the Cloudformation template, so that based on the relevant region, the relevant resource can be spun up

Answer: C

Explanation: The optional "Mappings" segment matches a key to a set of named values. For example, if you want to set a region-based value, you can create a mapping using the name of the region as a key and containing values for each particular region. You use the intrinsic function Fn::FindInMap to get values from a map.

183. The company for which you are working has an enormous infrastructure built on AWS. However, there are some security concerns regarding this infrastructure and an external auditor has been given the task of thoroughly checking all the AWS assets of your company. Your company is located in the Asia-Pacific (Sydney) region of AWS whereas the auditor is in the USA. You have been assigned the task of providing the auditor with login in order to check all your VPC assets, in particular, security groups and NACLs. Choose the best and secured solution for initiating this investigation.

A. Give him root access to your AWS Infrastructure; As he is an auditor, he will need access to every service

B. Create an IAM user who will have read-only access to your AWS VPC infrastructure and provide the auditor with those credentials

C. Create an IAM user tied to an administrator role. Also, provide an additional level of security with MFA

D. Create an IAM user with full VPC access but set a condition that will not allow him to modify anything if the request is from any IP other than his own

Answer: B

Explanation: It is only providing the required permissions as high level permissions should be avoided. So you can choose this option as it is the best fit for this requirement.

184. As an orchestration tool for the control of pipelines, AWS CodePipeline has been used in a financial firm. Certain AWS services, like CodeCommit, CodeBuild, CodeDeploy, etc, can be installed in pipelines. There is a safety policy in place to manage all artifacts created during the execution of the pipeline in a durable and highly available location. What is the correct statement about handling an objects by AWS CodePipeline? (Choose 2)

A. When CodeBuild is used in the build stage, the user can only configure input artifacts but not output artifacts

B. User can select a custom S3 bucket as the artifact store while creating a new pipeline
C. When artifacts are saved in an S3 bucket, the bucket can belong to a different region
D. Users can configure a default S3 bucket as the artifact store in the same region and account as the pipeline
E. An artifact can be stored in S3 or EC2 EBS

Answer: B and D

Explanation: The client may choose an S3 bucket default or a custom S3 bucket to store artifacts when building pipelines. By default, for each pipeline, a dedicated folder is created in the s3 bucket.

185. You have designed a critical application with the requirement of the least downtime of rollback (if required). You want to rollout updates to your application introduced on Elastic Beanstalk Environment. Which of the following deployment action is best suitable for this purpose?

A. Create another parallel environment in Elastic Beanstalk. Use the swap URL feature

B. Create a CloudFormation template with the same resources as those in the Elastic Beanstalk environment
C. Create another parallel environment in Elastic Beanstalk. Create a new Route53 Domain name for the new environment and release that URL to the users

D. Use Rolling updates in Elastic Beanstalk so that if the deployment fails, the rolling update feature would roll back to the last deployment

Answer: A

Explanation: As the requirement is the least downtime, it is ideal for creating a blue green deployment environment and use the Swap URL function to swap new deployment environments and then swap back in case of deployment failure.

When you update your application versions, Elastic Beanstalk implements the in-place update, which may make your application unavailable for a short period of time.

This downtime can be prevented by a blue-green deployment where the new version is deployed in a different environment. CNAMEs of both environments can be swapped, and the traffic can be redirected to a new version instantly.

186. The IT department of a company wants to launch instances in the Auto-scaling group. They need to setup lifecycle hooks for setting custom based software and do the required configuration on the instances. This setting would take an hour at maximum. Considering this scenario, how will you suggest to setup lifecycle hooks? Select 2 answers.

A. Configure the lifecycle hook to record heartbeats. If the hour is up, restart the timeout period
B. Configure the lifecycle hook to record heartbeats. If the hour is up, choose to terminate the current instance and start a new one
C. If the software installation and the configuration is complete, then send the signal to complete the launch of the instance
D. If the software installation and the configuration is complete, then restart the time period

Answer: A and C

Explanation: The instance will wait for an hour by default, and Auto-scaling will carry on with the launch or terminate the process (Pending:Proceed or Terminating:Proceed). You can reboot the timeout

period by recording a heartbeat if you need more time. You can finish the lifecycle action that continues the launch or termination process if you finish before the time limit ends.

187. Your company has announced a new IT procedure for EC2 instances, which states that EC2 instances must be of the particular instance type. You want to find out the list of instances that do not match the demanded instance type. From the following options, which one would you use to obtain the list?

A. Use TrustedAdvisor to check which EC2 instances do not match the intended instance type

B. Use VPC Flow Logs to check which EC2 instances do not match the intended instance type
C. Use AWS CloudWatch alarms to check, which EC2 instances do not match the intended instance type
D. Use AWS Config to create a rule to check EC2 instance type

Answer: D
Explanation: You can create a rule in AWS Config that can check whether EC2 Instances follow a particular type of instance.

188. Your company IT supervisor is interested in optimizing the cost of running AWS resources. Which of the 2 options are suitable for this purpose?

A. Use the Trusted advisor to see the underutilized resources

B. Create a script, which monitors all the running resources and calculates the cost accordingly. It analyses those resources and sees which can be optimized
C. Create CloudWatch alarms to monitor underutilized resources and either shutdown or terminate resources, which are not required
D. Create CloudWatch logs to monitor underutilized resources and either shutdown or terminate resources, which are not required

Answer: A and C

Explanation: CloudWatch alarms can be used to see if the resources threshold level, for a long time period, is below or not. If it is below the threshold level, then you can decide whether to stop or terminate the resources.

With Trusted Advisor, you will obtain all kinds of checks that can be used to optimize or reduce the costs of your AWS resources when you enable the Cost Optimization section.

189. Jenifer has an Auto-scaling group with the following setting: Minimum capacity: 2, Desired capacity: 2 and Maximum capacity: 4. The launch setup has AMIs that are t2.micro-instance-based. The program that runs on these instances now faces problems and she has found that the solution is to change the type of instance in the Auto-scaling group. From the following, which meets the desired requirement?

A. Change the desired and maximum size of the Auto-scaling group to 4. Make a copy of the launch configuration. Change the instance type in the new launch configuration. Attach that to the Auto-scaling group. Change the maximum and Desired size of the Auto-scaling Group to 2

B. Make a copy of the Launch configuration. Change the instance type in the new launch configuration. Attach that to the Auto-scaling Group. Change the maximum and Desired size of the Auto-scaling Group to 4. Once the new instances are launched, change the Desired and maximum size back to 2

C. Delete the current Launch configuration. Create a new launch configuration with the new instance type and add it to the Auto-scaling Group. This will then launch new instances

D. Change the Instance type in the current launch configuration. Change the Desired value of the Auto-scaling Group to 4. Ensure the new instances are launched

Answer: B

Explanation: The launch configuration must be copied and a new instance type added. The Auto-scaling group changes to include the new type of instance. Switch to value 4 as is the desired number in the Auto-scaling

group so that new instance instances can be started. Once it is released, switch the requested size back to 2 to delete the instances with the older setup. Please note that the current instances are equally distributed over several AZs, as Auto-scaling uses the AZRebalance method as its first means of terminating instances.

190. You are running a video processing application launched in AWS. Users upload videos on site, which are then processed by using the built-in custom program in case if there are any failures in processing, the program will be able to balance the situation. Now considering the minimum cost budget, which of the following mechanism should you use to deploy the instance for running video processing activities?

A. Create a launch configuration with Spot Instances. Ensure the user section data details the installation of the custom software. Create an Auto-scaling group with the launch configuration
B. Create a launch configuration with Dedicated Instances. Ensure the user section data details the installation of the custom software. Create an Auto-scaling group with the launch configuration
C. Create a launch configuration with On-Demand Instances. Ensure the user section data details the installation of the custom software. Create an Auto-scaling group with the launch configuration
D. Create a launch configuration with Reserved Instances. Ensure the user section data details the installation of the custom software. Create an Auto-scaling group with the launch configuration

Answer: A

Explanation: The application should be able to recover the failures and solutions should be cost effective therefore, spot instances are best for this purpose. The launch configuration can be used to request spot instances.

191. You are an AWS DevOps engineer in a company and have created a job in Jenkins to use an Elastic Beanstalk script. The eb config command was used to set up an Elastic Load Balancer

Security Group. After a while, one of your colleagues thought that the security group should be changed to another and modified the configuration option in the .ebextensions folder with the config file. The Security Group, however, did not change to the new Jenkins job when it was rebuilt. What is the issue and how are you going to resolve it?

A. The settings in .ebextensions folder cannot override that in the EB CLI command. Remove the configuration settings for the Security Group in eb config command

B. When there is a conflict for the configurations in .ebextensions and eb config, the default value of the option is used. The Security Group settings in eb config should be removed to avoid a conflict
C. The .ebextensions config file was processed first. The eb config ran next and overrode the .ebextensions config file. Modify the script in Jenkins so that the .ebextensions config file is processed after eb config
D. Security Group name cannot be changed in .ebextensions config file. The eb config in the script needs to be modified to use the correct Security Group

Answer: A

Explanation: The key to this question is that when Elastic Beanstalk settings are applied, there are different priorities. The AWS Management Console, EB CLI, AWS CLI and SDK settings are the safest and most appropriate settings for those directly applicable to the system. So, Option A is the best suit for the given question.

192. On the SaaS platform, an email with important information is sent to the user when an existing user modifies his/her subscription information. The subscription data is stored in a DynamoDB table and the stream allows the changes to table item to be recorded. The stream data is tracked and emails sent accordingly through a Java program. Which statement is correct about the use of DynamoDB streams?

A. The information in the Stream log is stored forever

B. DynamoDB Streams can only write a stream record when items in the table are created or updated
C. The data in DynamoDB Streams is also encrypted
D. The Java application can only view the data items as they appeared after they were modified

Answer: C
Explanation: The data in both DynamoDB table (at rest) and DynamoDB Streams is fully encrypted. All other options are not correct.

193. There are many onsite servers operated by a company who plans to move its main servers to AWS ap-southeast-2 in order to formulate a disaster recovery plan in the cloud. The servers are installed on Microsoft HyperV-managed virtual machines for Windows. The company's team wants to use an AWS platform to ease the migration process by periodically creating AMIs. The tool can also schedule replications and monitor progress. Which tool should be used to fulfill the requirement?

A. AWS Database Migration Service

B. AWS DataSync
C. AWS Server Migration Service
D. AWS Migration Hub

Answer: C
Explanation: AWS Server Migration Service is ideal to simplify the AWS migration by setting up AMI ready to deploy on Amazon EC2 for on-site virtual machines within Microsoft Hyper-V or SCVMM.

194. The technical team of your company is concerned with AWS account security. What will you suggest to prevent the account from being hacked?

A. Do not write down or remember the root account password after creating the AWS account

B. Use a short but complex password on the root account and any administrators
C. Use AWS IAM Geo-Lock and disallow anyone from logging in except for in your city
D. Use MFA on all users and accounts, especially on the root account

Answer: D

Explanation: The user can add an additional protective layer on top of the username and password using MFA. When MFA is enabled, the user is prompted for a username and password (first factor, what user knows), and authentication code of the AWS MFA system (second factor, what user has) to be registered on an AWS website.

195. Your company's owner asked you to show all of the CloudFormation stacks, which have a completed status. Which command should you use?

A. list-stacks
B. stacks-complete
C. list-templates
D. describe-stacks

Answer: A

Explanation: This command returns the stack information of all the stacks whose status is matched with StackStatusFilter. If the command does not find any stack filter, it will return the information of all stacks. Information of the deleted stacks is kept in record for 90 days.

196. A blue/green deployment approach for a new application is being applied by the Allen. A CloudFormation stack with tools like Auto-scaling group, launch setup and Classic Load Balancer allows the application to be deployed for every new release. Route53 configures a weighted routing system. For route53, a small percentage of traffic enters the green environment and the weight increases until the full production traffic is borne by the green environment. What is the downside of this approach?

A. Only a classic load balancer is suitable for this approach. A network load balancer or application load balancer cannot be used together with the Route53 weighted routing policy

B. Route53 weighted routing policy is not suitable if a rollback is required
C. DNS TTL decides how long clients cache query results. Certain sessions may still be tied to the previous environment so that it may impact the speed to deploy a new environment or rollback
D. Weighted routing policy is complicated to implement in a pipeline

Answer: C

Explanation: One of the downsides of blue/green deployment is that older clients may take longer timers until they use the new environment. DNS TTL decides how long the cache requests output for the cache users.

Nonetheless, some sessions may still be related to the previous setting of older clients and potentially misbehaving clients in the wild.

For further detail visit:

https://d1.awsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf

197. Leo has just started using AWS ECS/ECR for its Docker applications. He is searching for a pipeline system that can use a blue/green configuration to deploy the container software in the ECS cluster. At present, all Docker pictures are stored as artifactories in the ECR system in pipelines. How is this pipeline to be applied the quickest?

A. In AWS CodeDeploy, create a CodeDeploy application and a deployment group to configure the blue/green deployment to ECS cluster using the image in the ECR repository

B. In AWS CodePipeline, add a source stage for ECR docker image and a deployment stage for ECS where the deployment runs with a CodeDeploy application and deployment group
C. In AWS CodePipeline, configure a source stage for ECR, a build stage with CodeBuild for the docker image and a

deployment stage using CloudFormation to deploy ECS cluster

D. Create a Jenkins server in EC2 instance. In the Jenkins job, add a source stage for ECR docker image and a deployment stage for the ECS cluster using AWS CLI commands

Answer: B

Explanation: In CodePipeline, two-stage services will fulfill the necessity. The action provider could be selected as Amazon ECS (Blue / Green) during its deployment stage. The use of the Jenkin server is not good for this requirement. Using CloudFormation for deployment of Blue/Green deployment needs more effort as compared to the use of CodePipeline because the pipeline cannot be complete in AWS CodeDeploy. The pipeline needs an origin step, for instance, if the new image is uploaded to Amazon ECR, it can deploy container images automatically.

198. What is the meaning of the given code in the CloudFormation template?

```
“SNSTopic” : {
        “Type” : “AWS::SNS::Topic”,
        “Properties” : {
        “Subscription” : [{
              “Protocol” : “sqs”,
        “Endpoint” : { “Fn::GetAtt” : [ “SQSQueue” , “Arn” ] }
           }]
        }
```

A. It creates an SNS topic and adds a subscription ARN endpoint for the SQS resource created under the logical name SQSQueue

B. It creates an SNS topic and then invokes the call to create an SQS queue with a logical resource name of SQSQueue

C. It creates an SNS topic that allows SQS subscription endpoints

D. It creates an SNS topic, which allows SQS subscription endpoints to be added as a parameter on the template

Answer: A

Explanation: The function Fn::GetAtt returns the value of an attribute from any resource in the template.

199. A large number of aerial image data has been uploaded to S3 by your company. In the past, you used a dedicated group of servers in your local environment to process these data and used Rabbit MQ– an open source message system to get job information to the servers. The data would go to the tape and be shipped offsite once processed. Now your manager told you to use the current design along with AWS archive storage and messaging to reduce costs. Which option is right?

A. Setup Auto-scale workers triggered by queue depth that use spot instances to process messages in SQS. Once data is processed, change the storage class of the S3 objects to Glacier

B. Use SNS to pass job messages and use CloudWatch alarms to terminate spot worker instances when they become idle. Once data is processed, change the storage class of the S3 object to Glacier

C. Use SQS for passing job messages. Use CloudWatch alarms to terminate EC2 worker instances when they become idle. Once data is processed, change the storage class of the S3 objects to Reduced Redundancy Storage

D. Change the storage class of the S3 objects to Reduced Redundancy Storage. Setup Auto-scaled workers triggered by queue depth that use spot instances to process messages in SQS. Once data is processed, change the storage class of the S3 objects to Glacier

Answer: A

Explanation: RabbitMQ was used internally as a messaging service that is why SQS should be used. Therefore option B is invalid as it is using SNS.

The best option for reducing costs is Glacier, as everything was stored on tape at the on- site location. So, Option C is therefore out as it is using RRS. Option D is not valid as there is no need to use the storage class RRS to put the file in it.

Hence option A is more suitable as it is using Glacier for processed data.

200. A company is using an Auto-scaling group to scale out and scale in EC2 instances. The traffic peak occurs every Monday at 8 am and it comes down before the weekend on Friday at 5 pm. If you have to configure Auto-scaling group in this scenario, what would you do?

A. Create a scheduled policy to scale up on Friday and scale down on Monday

B. Create a scheduled policy to scale up on Monday and scale down on Friday
C. Create dynamic scaling policies to scale up on Monday and scale down on Friday
D. Manually add instances in the Auto-scaling group on Monday and remove them on Friday

Answer: B
Explanation: Scheduling based scaling allows you to scale according to the change in a load of the application. I.e., if the traffic level is high, then it will scale up, and when it is low, it will scale down.

201. What is the possible cost-effective solution to storing a large volume of data that is accessible for a short period and archived indefinitely after that?

A. Keeping all your data in S3 since this is durable storage

B. Storing your data in an EBS volume, and using lifecycle policies to archive to Amazon Glacier
C. Storing your data in Amazon S3, and using lifecycle policies to archive to S3-infrequently-access
D. Storing your data in Amazon S3, and using lifecycle policies to archive to Amazon Glacier

Answer: D
Explanation: Amazon Glacier is a long term secure and durable storage service that is used for data archiving and long-term back-ups.
The configuration is a set of rules to define the action on the Amazon S3

bucket for a group of objects. Lifecycle configuration enables you to manage the S3 lifecycle. Following actions can be defined for Amazon S3:

Transition Actions: Transition actions describe the transition of one storage class to another storage class, for example, you want to change your S3 storage to Glacier after 30 days.

Expiration Actions: This action defines the expiration timeline of objects. S3 deletes the expired objects automatically.

202. Barrett just created a regional cluster Aurora MySQL in ap-southeast-2 on AWS console as part of a new venture. The Aurora cluster is highly available and durable, with a number of replicas in the Aurora database. What statement is correct about the replicas of the Aurora database? (Choose 2)

A. Replicas are configured with smaller storage volume than the primary DB instance to save cost

B. By default, there is only one replica generated in each availability zone
C. The replicas are located in different availability zones in the same region as the primary database, which is ap-southeast-2
D. The replicas can share both the read and write traffic
E. Aurora automatically fails over to an Aurora Replica if the primary DB instance becomes unavailable

Answer: C and E

Explanation: A primary DB instance and Aurora replicas are included in an Amazon Aurora Database cluster. Since replicas are distributed over multiple availability zones to ensure that even if one availability region is out of operation, the server still operates. In the case of primary DB failover to an Aurora, replica is the key advantage of Aurora.

203. Using ELB, Auto-scaling group of Java/Tomcat application servers, and DynamoDB as a data store in EC2 instance, a web-startup is running its successful social news app. Web applications require high memory therefore, m2x large is the most suitable instance. The semi-automated creation and testing of a new AMI for the application servers is required by each new deployment, which

takes some time and is done only once a week. A new chat feature was recently introduced in nodejs and is waiting to be integrated into the architecture. The new component is shown CPU bound in the first test because the company has some experience using Chef, and has decided to streamline the deployment process and use AWS OpsWorks as an application lifecycle tool to simplify application management and reduce deployment cycles. What configuration is needed in AWS OpsWorks to integrate the new chat module into the most cost-effective solution?

A. Create one AWS OpsWorks stack, create two AWS OpsWorks layers and create one custom recipe
B. Create two AWS OpsWorks stacks, create two AWS OpsWorks layers and create one custom recipe
C. Create two AWS OpsWorks stacks, create two AWS OpsWorks layers and create two custom recipe
D. Create one AWS OpsWorks stack, create one AWS OpsWorks layer and create one custom recipe

Answer: A

Explanation: Only one Opswork stack with two layers can be used, one for Node.js and one for the standard app.

The configuration of your entire application is defined by an AWS OpsWorks Stack: load balancers, server software, database, etc. You control each part of the stack by building layers that define the software packages and other configuration details such as Elastic IPs and security groups. You can also deploy your software on layers by identifying the repository and using Chef Recipes optionally to automate everything Chef can do, such as creating directories and users, setting up databases, etc. You can use the built-in automation of OpsWorks Stacks to scale your application and recover from instance failures automatically. You can check and control who can view and manage the resources your application uses, including SSH access to the instances your application uses.

204. You are running an application globally. You have multiple EC2 instances running in different regions. You want to monitor the performance of each instance using CloudWatch. What will you do?

A. Create a separate dashboard in each region

B. This is not possible
C. Register instances running on a different region to CloudWatch
D. Have one single dashboard to report metrics to CloudWatch from different regions

Answer: D

Explanation: AWS resources can be monitored through a single CloudWatch dashboard in several regions. For example, you can create a dashboard that displays the use of CPUs for an EC2 instance in the us-west-2 region with your billing metrics in the us-east-1 region.

205. AWS Codepipeline is used by Bert as a tool of continuous integration and deployment for a Lambda function. In its deployment stage, the deployment provider is CloudFormation and the Action Mode is configured as “Create or update a stack”. Initially, the pipeline was working fine but there has been an inappropriate update in CloudFormation stack due to some mistake in the template. Now, Bert is asked to create a design in which updating directly on stack is not done and instead, there must be a way to preview the stack changes before the deployment is completed. How should he work on this task?

A. Modify the Action Mode to “Create or replace a change set” for users to review the change. Then add another deployment stage to execute the change set if the change set is approved

B. Change the Action Mode to “Delete a stack” so that the environment is clean. Then add another stage to deploy the CloudFormation stack using the new template
C. Modify the deployment provider to CodeDeploy to avoid any issues on the CloudFormation stack update
D. Change the code review process so that any issues on CloudFormation stacks are avoided

Answer: A

Explanation: The change set is created for analyzing the changes in the

resources before execution, by using “Create or replace a change set” action mode. If changes appear correct, the change set is executed.

206. An online web store adopted AWS CloudFormation to automate load-testing of the details of their online products. They created two templates of CloudFormation, one is for the details of their products and the other one for load testing stack. Load-testing stack creates an RDS Postgres database and two web servers running on EC2 instance that measures response time, sends HTTP requests, and records the results into the database. Test time is usually 15 – 30 minutes. The AWS CloudFormation stacks are torn down immediately after the test completion. The recorded test results of Amazon RDS database must remain accessible for virtualization and analysis.

If the AWS CloudFormation load-testing stack is deleted, then what could be the possible solutions that allow access to the test results. (Choose 2)

A. Define an updated policy to prevent the deletion of the Amazon RDS database after the AWS CloudFormation stack is deleted

B. Define a DeletionPolicy of type Retain for the Amazon RDS resource to assure that the RDS database is not deleted with the AWS CloudFormation stack
C. Define a DeletionPolicy of type Snapshot for the Amazon RDS resource to assure that the RDS database can be restored after the AWS CloudFormation stack is deleted
D. Define automated backups with a backup retention period of 30 days for the Amazon RDS database and perform point-in-time recovery of the database after the AWS CloudFormation stack is deleted
E. Define an Amazon RDS Read-Replica in the load-testing AWS CloudFormation stack and define a dependency relation between master and replica via the Depends On attribute

Answer: B and C

Explanation: If the user wants to preserve any resource, then DeletionPolicy should be attached to it. DeletetionPolicy attributes allow the user to preserve or in some cases backup a resource when its stack is deleted. AWS CloudFormation deletes a resource by default if no DeletionPolicy attribute is

attached to it.
If the user specifies the Retain for any source, this will prevent the resource deletion if its stack is deleted.

207. When considering AWS Elastic Beanstalk, the 'Swap Environment URLs' is helpful in which of the following deployments?

A. Immutable Rolling Deployments

B. Blue-Green Deployments
C. Canary Deployments
D. Mutable Rolling Deployments

Answer: B
Explanation: When you update your application versions, Elastic Beanstalk implements the in-place update, which may make your application unavailable for a short period of time. This downtime can be prevented by a blue- green deployment where the new version is deployed in a different environment. CNAMEs of both environments can be swapped, and the traffic can be redirected to a new version instantly.
Deployments in blue/green require your environment to run independently of your production database if your application uses one. If you have an Amazon RDS DB instance attached to your environment, the data will not be transferred to your second environment and will be lost if the original environment is terminated.

208. A team uses many Lambda functions to develop micro-services. The group has consulted Aliana on how the project can be designed. The pipeline needs to be AWS-CodePipeline, which enables the Lambda code to be continuously generated and the applications deployed in AWS, wherever a team member has a new Git commit. What service combinations are appropriate for that scenario in AWS CodePipeline? (Choose 2)

A. Source stage (CodeCommit), Build stage (Travis CI), Deployment stage (CloudFormation)

B. Source stage (Github), Build stage (Jenkins), Deployment stage (CodeDeploy)
C. Source stage (Bitbucket), Build stage (CodeBuild), Deployment stage (Amazon ECS)
D. Source stage (GitHub), Build stage (Jenkins), Deployment stage (Elastic Beanstalk)
E. Source stage (CodeCommit), Build stage (CodeBuild), Deployment stage (CloudFormation)

Answer: B and E

Explanation: Different services within CodePipeline can be implemented at various stages, which is highly flexible for users. Nevertheless, not every service is suitable in this particular case, even if it may be valid for another scenario. When a new pipeline is created, then you need to configure the source, build and deploy stages:

- Source stage include- CodeCommit, ECR, GitHub, and S3
- Build stage include- Jenkins or CodeBuild
- Deploy stage includes- CodeDeploy, ElasticBeanstalk, CloudFormation, S3, ECS and Service Catalog

Option D is close to the answer but Elastic Beanstalk is used for deployment of EC2 instance not for Lambda. While all other options are invalid.

209. Which system architecture is best for a company that is working on automatic photograph tagging by using Artificial Neural Networks (ANNs), which have C++ format and its processes on GPU? The images loaded in the S3 bucket are in millions but on average, 3 images per day. You control the S3 bucket for you in a batch. You have control on one more S3 bucket, in which a customer publishes JSON formatted manifest. Bootstrap time of your neural network software is 5 minutes, and an image takes 10 milliseconds to process using full GPU. Tags are JSON formatted, which must be published to S3 bucket.

A. Make an S3 notification configuration, which publishes to AWS Lambda of the manifest bucket. Make the Lambda CloudFormation stack, which contains the logic to construct an Auto-scaling worker tier EC2 G2 instances with the artificial

neural network code on each instance. Handle the CloudFormation Stacks creation success or failure using another Lambda function. Create an SQS queue of the images in the manifest. Tear the stack down when the queue is empty

B. Deploy your artificial neural network code to AWS Lambda as a bundled binary for the C++ extension. Make an S3 notification configuration on the manifest, which publishes to another AWS Lambda running controller code. This controller code publishes all the images in the manifest to AWS Kinesis. Your ANN code Lambda function uses the Kinesis as an Event source. The system automatically scales when the stream contains the images
C. Create an OpsWorks stack with two layers. The first contains lifecycle scripts for launching and bootstrapping an HTTP API on G2 instances for image processing, and the second has an always-on instance, which monitors the S3 manifest bucket for new files. When a new file is detected, requests instances to boot on the new artificial neural network layer. When the instances are booted and the HTTP APIs are up, submit processing requests to individual instances
D. Create an Auto-scaling, Load Balanced Elastic Beanstalk worker tier Application, and Environment. Deploy the artificial neural network code to G2 instances in this tier. Set the desired capacity to 1. Make the code periodically check S3 for new manifests. When a new manifest is detected, push all of the images in the manifest into the SQS queue associated with the Elastic Beanstalk worker tier

Answer: A

Explanation: The S3 Events are the best way to be informed when the images are sent to the bucket. You do not need to provide infrastructure here in advance, and since the S3 source provides event management, this should be used.

Amazon S3 can publish events (e.g., when an object is created in a bucket) to AWS Lambda and use your Lambda function as a parameter by passing the event data. This integration allows you to write Amazon S3 events Lambda

functions. In Amazon S3, you add bucket notification settings that identify the type of event you want Amazon S3 to publish and the Lambda function you want to invoke.

Further information as to why the second function of Lambda is required:

You can use AWS Lambda to create a CloudFormation stack. CloudFormation stack creation is an asynchronous call, so we do not have to wait until the whole stack moves to FAILED/SUCCEEDED state. In the CloudFormation advance section, you can get the notification of the status of the stack via SNS notification.

210. The organization Byron works for is creating an app and he must decide which AWS services to use. Customers should be able to upload new pictures to an S3 bucket via this Mobile App. A Lambda function is triggered when a new object has been uploaded. There is a couple of processing steps later on. The Lambda function, for example, calls the Recognition API and Amazon Recognition service. Metadata including both size and format is extracted in another Lambda function. The data is ultimately stored in a DynamoDB. What AWS software would he use to manage image analysis tasks?

A. Use a Jenkins pipeline to add several steps, which manage the various tasks for this application

B. Create an AWS Batch Job to manage and coordinate the tasks for this application
C. Use a CloudFormation template to create several SQS queues and Lambda functions to manage the workflow for this application
D. Use AWS Step Functions to manage the steps and tasks into a serverless workflow

Answer: D

Explanation: In this scenario, a number of steps need to be organized by an AWS system. As the AWS Step Functions are suitable to handle a number of services, the serverless workflow will turn the workflow into a state machine. The most useful is the AWS Step Functions. If the query states that several steps are necessary, AWS Step functions are the first functionality to be considered.

211. An AWS organization with several OUs (Organizational Units) is formed in a major company. The security administrator needs to create several CloudWatch Events rules due to several new security criteria. For policy, it is necessary that an SNS alert is triggered when certain AWS services are unexpectedly used that were not used by an OU for 6 months prior to this. How would you receive the latest system information to identify the possible services that can be included in the CloudWatch Event rule?

A. In AWS Resource Groups service, add a resource group for the candidate AWS services. View the last-accessed information for the group

B. In the IAM access advisor, view the service-last-accessed information for each OU to help identify which services should be added in the CloudWatch Event rule
C. Send an AWS CLI command to AWS Config service to get the service-last-accessed information for each organizational unit
D. In AWS Organization console, click each OU and view the service-last-accessed information

Answer: B

Explanation: IAM access advisor is able to find the most last-accessed data for every OU.

212. Albert is using AWS CodeBuild service to handle the build task in a CI/CD pipeline. In the pre-build phase of buildspec.yml, there is a docker login command such as “docker login –u $USER_NAME –p $LOGIN_PASSWORD”. And its user name and password are provided as variables in the env phase in the same buildspec.yml file. The credentials may be exposed, which is one of the security issues. What is the best way to overcome this issue?

A. Store the credentials in a file and put the file in an S3 bucket. Encrypt the S3 bucket via SSE-S3. Modify the buildspec.yml file to use the encrypted file in the S3 bucket

B. In the env phase of the buildspec.yml file, use the parameter-

store to specify the user name and password. The values are stored in Systems Manager parameter store

C. Store the buildspec.yml file in AWS CodeCommit rather than GitHub as IAM rules can be configured in CodeCommit to ensure the security
D. Add a strong IAM rule in AWS CodeBuild to make sure that only limited users can access the buildspec.yml file

Answer: B

Explanation: Systems Manager Parameter store is an ideal place to store the confidential data. In the buildspec.yml file, you can use the parameter store in env phase. In form of key-value pair.

213. Harrison plans to move his on-premises server to save the cost and for the disaster recovery plan. In order to configure the servers properly, he plans to use AWS SMS. From the following option, which is not a good use case for AWS SMS?

A. A data center running in Microsoft System Center Virtual Machine Manager

B. A Windows10 Pro server running in Microsoft Hyper-V
C. An Ubuntu16.04 linux server running in VirtualBox
D. A Java application running in a VMware vSphere Enterprise machine

Answer: C

Explanation: For AWS SMS, Oracle VirtualBox is not supported. Because servers in VMware vSphere, Microsoft Hyper-V / SCVMM, and Azure virtual machines can use Server Migration Service to duplicate server VMs as AWS AMIs, then EC2 instances can be generated using the AMIs.

214. A company has given you the task to configure an AWS Elastic Beanstalk work tier for easy debugging, but you are facing problems in finishing queue jobs, what should you configure?

A. Enhanced Health Reporting

B. Blue-Green Deployments

C. Rolling Deployments
D. Dead Letter Queue

Answer: D

Explanation: Elastic Beanstalk worker environment supports Amazon SQS queue service dead letter queues. In the dead letter queue, other queues can send messages that for some reasons could not be processed. Messages that are unsuccessful in the processing are targeted from source queue to the dead-letter queue. You can gather these types of messages in dead-letter queues to find the reason for their failure.

215. An IT company has various applications with end users around the globe. Route53 is used to route the traffic. Specific implementation techniques are applied, depending on the different application features. Like, Canary is used by some sensitive applications, while Blue/Green is used by others. Many programs also use Route53 for disaster recovery plan implementation. In order to carry out these strategies, what Route53 routing policy is required?

A. Weighted for Blue/Green. Weighted for Canary. Failover for Disaster recovery

B. Weighted for Blue/Green. Latency for Canary. Failover for Disaster recovery
C. Simple for Blue/Green. Weighted for Canary. Latency for Disaster recovery
D. Latency for Blue/Green. Simple for Canary. Weighted for Disaster recovery

Answer: A

Explanation: Weighted can be used both for Blue/Green and Canary deployments as this is an easy way of routing traffic to two stacks as appropriate. Failure to the route can assist in configuring an active-passive failover that can be used as a disaster recovery technique.

216. In order to construct a pipeline for the Docker image and to transfer it to the AWS ECR, a DevOps engineer is installed in AWS CodePipeline. The codepipeline configures the source stage to be

AWS CodeCommit for the source data. The service provider is AWS CodeBuild in the development phase. There are three stage; pre-design, build, and post construct stages for its buildspec.yml folder. It is the building stage in the pipeline, which generates the picture and transfers it to the ECR repository. What command should be put in the buildspec.yml file in the pre-build stage?

A. docker push $REPOSITORY_URI:latest

B. docker tag $REPOSITORY_URI:latest
C. aws ecr get-login --region $AWS_DEFAULT_REGION
D. docker build -t $REPOSITORY_URI:latest

Answer: C

Explanation: In the buildspec.yml file, there is collection of build commands and related configuration to run build for AWS CodeBuild. Preparation steps such as ecr login, the ECR database URI and tags definition should be used in pre-build phase.

217. Jack plans to launch instances that have an application installed with Auto-scaling. Which of the following methods will help to ensure that the instances are up and running for traffic from users in the shortest possible time?

A. Use user data to launch scripts to install the software

B. Log in to each instance and install the software
C. Use AMIs, which already have the software installed
D. Use a Docker container to launch the software

Answer: C

Explanation: As you are using the AMI, which already has the required software installed, it will therefore implement the fastest way to launch an instance. You can configure the public AMI as a custom AMI with your own defined configuration. The instance launched from that AMI will contain all the modifications you have made.

218. Arthur is working in a company that has recently extended its data center into a VPC on AWS. The on-premises users are required

to manage AWS resources from the AWS console. He is restricted from re-creating IAM users. Which of the options below will fit his authentication needs?

A. Use on-premises SAML 2 O-compliant identity provider (IDP) to grant the members federated access to the AWS Management Console via the AWS single-sign-on (SSO) endpoint

B. Use on-premises SAML2.0-compliant Identity Provider (IDP) to retrieve temporary security credentials to enable members to sign in to the AWS Management Console
C. Use Auth 2.0 to retrieve temporary AWS security credentials to enable members to sign in to the AWS Management Console
D. Use on-premises SAML2.0-compliant Identity Provider (IDP) to retrieve temporary security credentials to enable members to sign in to the AWS management console

Answer: A

Explanation: He can use a role to set up your SAML 2.0 IDP and AWS so that federated users can access the AWS management console. The role empowers the user to perform tasks in the console.

219. Charlie is working for an organization that has an on-premises infrastructure. There is a plan to move to AWS. The arrangement is to move the development environment first. There is a number of custom based applications that should be deployed for the development community. Which of the following can actualize the application for the development group? (Choose 2)

A. Create Docker containers for the custom application components

B. Use OpsWorks to deploy the docker containers
C. Use CloudFormation to deploy the docker containers
D. Use Elastic Beanstalk to deploy the docker containers

Answer: A and D

Explanation: Elastic Beanstalk allows the use of Docker containers for web application deployment. You can create your own runtime environment using Docker containers. You can choose your own framework, programming language and any software dependencies not provided by any other platform (for example package managers or tools). Dockers include all configuration information and applications that your web application needs to be running is automatically included.

220. Mark works in a company that uses AWS tools. One of the main safety measures is ensuring encryption of all information both in rest and in transit. Which one is the correct implementation in line with this policy?

A. Enabling sticky sessions on your load balancer

B. Enabling Proxy Protocol
C. Using Server-side encryption for data encryption in transit and SSL termination on ELB for data encryption at rest
D. Using S3 Server-side encryption for data encryption at rest and SSL termination on ELB for data encryption in transit

Answer: D
Explanation: Enabling S3 SSE encryption will encrypt data at rest for EBS volumes and SSL termination allows encrypted traffic between the client and ELB. If SSL termination is not enabled, then there is a need to use layer 4 for encryption that is not supported by sticky session.
Proxy protocol functionality gives you additional insight into your clients' communication data by using TCP load balancing.

221. You support the development team to create a new web application with an Aurora MySQL cluster. The cluster was built with a Dev/Test model as the specifications were not clear. The DB cluster had only one writer and the read replica was not available. The Aurora database has recently proven successful and you have been directed to configure the DB cluster to make it available. What should you do to set up the current cluster Aurora?

A. Delete the cluster and create a new Aurora cluster with

serverless features to make it highly available

B. In RDS console, add Aurora Replicas to the DB cluster by actions -> add reader
C. Since an active Aurora cluster cannot configure read replicas, delete and recreate the Aurora cluster with a Production template
D. Use AWS CLI to configure cross region replications for the DB cluster to make it highly available

Answer: B

Explanation: In order to make Aurora highly available, you can add Aurora replicas by using the RDS console to the DB cluster. Removing the cluster and recreating is not the proper solution, also cross region replication is not needed so this option is also incorrect. Serverless features may improve capacity but availability of both the serverless or provisioned cluster is the same.

222. Martin is building a CloudFormation stack using an AWS CodePipeline. The stack is used in a human approval project to deploy resources that allow the implementation of a state machine to pause for a task. Once the client has approved the project, the process begins. CloudFormation already has resources including the Lambda function that sends an email link for approval, the API Gateway endpoint, which triggers Lambda function, and an SNS subject that emails for approval. What is another important resource to build in the template?

A. An SWF resource in the CloudFormation type AWS::SWFFunctions::Activity

B. A state machine resource created in the type of AWS::StepFunctions::Activity
C. An SWF state machine resource configured in AWS::SWFFunctions::StateMachine
D. A Step Function resource created in the type AWS::StepFunctions::StateMachine

Answer: D

Explanation: In this query, after the client has approved the task the workflow progresses to the next level. It can be run using a state machine because the Step Function state machine as a key tool for obtaining the manual approval function during the process.

223. AWS CodePipeline has been designed to build and implement an AWS Step Feature. The AWS Step Function implements a state machine for querying a large number of DynamoDB records. Now, after a new version has already been introduced, your lead has asked you to add an additional test stage in the present pipeline. In the new phase, which service will you use to begin the test step?

A. Use a Lambda function to trigger the Step Function in order to do the testing

B. Configure an API endpoint in API Gateway. The API will call StartExecution to start the state machine of the Step Function
C. Add a CloudWatch Event rule as the action provider in the new stage. Configure the Step Function as the target of the Event rule
D. Add a Jenkins job to invoke the Step Function to perform any testing needed

Answer: A

Explanation: A custom stage with an action group can be added to CodePipeline. Then the user can adjust to the requirements of the action group. There are several methods for the AWS Step Function to trigger the state machine. During the test phase, the step function of a Lambda function can be configured to invoke Step Function. To start Step Function, you can use CloudWatch events but it is not available as action type in Code Pipeline. Similarly, API Gateway is also not available. While the use of Jenkin will only create unnecessary complexity.

224. While working on the CloudFormation template, you want to use intrinsic functions to assign values to properties that will not be available until runtime. Choose the best description to use intrinsic functions.

A. You can use intrinsic functions only in the resource properties part of a template

B. You can use intrinsic functions in any part of a template, except AWS TemplateFormatVersion and Description
C. You can use intrinsic functions in any part of a template
D. You can only use intrinsic functions in a specific part of a template. You can use intrinsic functions in resource properties, metadata attributes, and update policy attributes

Answer: D

Explanation: AWS CloudFormation offers multiple built-in functions for the management of your stacks. To assign values to properties that are not available until runtime, use the intrinsic functions in your templates.

Only in certain parts of a template can intrinsic functions be used. Currently, intrinsic functions can be used to update policy attributes, outputs, Metadata and resource properties. You can also use intrinsic functions to build stack resources on condition.

225. A user has launched an EC2 instance using CloudFormation. He wants that Auto-scaling and ELB stack creation starts after the EC2 instance is launched and configured properly. What should be the possible way of configuration?

A. The user can use the WaitCondition resource to hold the creation of the other dependent resources

B. The user can use the HoldCondition resource to wait for the creation of the other dependent resources
C. It is not possible that the stack creation will wait until one service is created and launched
D. The user can use the DependentCondition resource to hold the creation of the other dependent resources

Answer: A

Explanation: You can use the WaitCondition to coordinate stack resource creation with external configuration actions and to track the configuration process status.

226. Max was given a mission to build a Jenkins job to do the backup for a very complicated system. The Jenkins job will activate the backup function every week. The backup itself consists of a number of steps including running preparations steps, launching DB snapshots, checking the status of DB snapshots, running post-scripts, etc. Every move can be made by a Lambda function. Therefore, the solution designer expects to see and monitor the operational status of an AWS-owned system for each phase. To achieve this, what service should Max use?

A. AWS EventBridge

B. AWS CodePipeline
C. AWS Step Function
D. AWS SNS

Answer: C

Explanation: In this problem, it is necessary to see every phase of the backup in an AWS system. Step functions may fulfill the need as they have a user interface to understand the running status of a state machine for every step. For e.g., like first to submit jobs, you get job status either completed or failed after a waiting period.

227. You were hired for a start-up company as a DevOps engineer. Your company uses AWS for 100% of its infrastructure. Currently, the deployment is not automated and has experienced many failures while trying to deploy to production. The company has told you that the risk mitigation process is the most important thing now, and you have enough funds for tools and AWS resources. Depending on the type, the company stack includes a 2-tier API with data stored in DynamoDB or S3. In Auto-scaling groups, the compute layer is EC2. company uses Route53 for DNS to an ELB. A load of an ELB balance over EC2 instances. The scaling group varies properly from 4 to 12 EC2 servers. Which of the following approaches, given the stack of this company and its priorities, best meets the needs of the company?

A. Model the stack in AWS OpsWorks as a single Stack, with 1

compute layer and its associated ELB. Use Chef and App Deployments to automate Rolling Deployment

B. Model the stack in three CloudFormation templates: Data layer, compute layer, and networking layer. Write stack deployment and integration testing automation following Blue-Green methodologies
C. Model the stack in AWS Elastic Beanstalk as a single Application with multiple Environments. Use Elastic Beanstalk's Rolling Deploy option to progressively roll out application code changes when promoting across environments
D. Model the stack in 1 CloudFormation template, to ensure consistency and dependency graph resolution. Write deployment and integration testing automation following Rolling Deployment methodologies

Answer: B

Explanation: You should use blue/green deployment and nested CloudFormation stack for deployment.

When infrastructure grows, common patterns may arise, in which each of your templates declares the same components. You can create dedicated templates and separate common components. In this way, you can mix and match various templates, but use nested stacks to create a single stack. To create other stacks with in a stack, one should use nested stacks. AWS::CloudFormation::Stack resource is used in your template to reference other templates to create nested stacks.

228. You work for a big company and the department is responsible for the new CI/CD pipeline to move from the premises to the AWS. The software uses a license that is limited to the number of vCPUs and the same license was agreed to be used on the AWS platform. The technology is built by Amazon AMI and several teams including QA and DEV are often using the Jenkins pipeline to launch new EC2 instances. Your team manager worries about the excess of the number of servers used and ask you to add a step in Jenkins to achieve the current license status consumption. What is the best way to do this?

A. Create a License Configuration for this license in AWS License

Manager. Call the CLI list-usage-for-license-configuration to get the license consumption status

B. Use a DynamoDB table to record the usage status of vCPU. Add a Jenkins step to read and write the table to manage the license consumption status
C. Create a Lambda to keep counting the number of servers and compare them with the license limit
D. Use a shell script to count the total number of EC2 instances in all teams and compare the number with the limit that the license allows

Answer: A

Explanation: This concern relates to the use of the license. AWS License Manager, which simplifies the process of getting software licenses to AWS, should be the first service to consider. License settings can be generated as vCPU with a license type: the settings can be associated with an AMI and the status of the use is then tracked.

229. A programmer asks to help set up a new Java project for AWS CodeBuild. Maven is responsible for the design. There are several steps in its buildspec.yml file such as install, pre-build and build. If a command fails in process, CodeBuild still has the opportunity to run some shell commands, such as log compilation, in order to carry out customs operations. How should the CodeBuild project be designed to meet the needs?

A. Trigger an SNS notification if a build fails in the CodeBuild project. Use a Lambda Function to subscribe to the SNS topic and handle the required custom operations

B. Create a CloudWatch Events rule for any CodeBuild failure event. Add a Lambda Function as the target to do custom operations
C. Put the custom operation commands in the post_build phase so that they can still run when the build fails
D. Add a finally block after the commands block in each phase. Put the custom shell commands in the finally block

Answer: D

Explanation: The buildspec.yml file has provided the final optional block for executing commands even if a command fails in the command block to overcome this situation.

230. An IT company's technical assistant comes to know that Elastic Beanstalk service provides a managed update facility, which is minor and patches version updates. The company starts hosting a production environment in Elastic Beanstalk. The technical assistant comes to you to ask the effects of an update on the system as the system requires these updates periodically. What would you tell him about managed update facility?

A. Elastic Beanstalk applies managed updates with no down time
B. Elastic Beanstalk applies managed updates with no reduction capacity
C. Package updates can be a configurable weekly maintenance window
D. All of the above

Answer: D

Explanation: In a configurable weekly maintenance window with Managed Platform Updates, you can set up your environment to apply minor and patch version updates automatically. Elastic Beanstalk applies managed updates without downtime or reduced capacity and immediately cancels the update if your application fails health checks of running instances when the new version executes your application.

231. A company is designing a CloudFormation template, which deploys a LAMP stack. Its users deployed the stack and CREATE_COMPLETE status is showing, but the apache server is still not up and running and is experiencing issues while starting. Now, the company wants that CREATE_COMPLETE status is shown only when all resources are completely up and running. What should be done to fulfill this requirement? (Choose 2)

A. Use lifecycle hooks to mark the completion of the creation and configuration of the underlying resource

B. Use the CFN helper scripts to signal once the resource configuration is complete
C. Define a stack policy, which defines that all underlying resources should be up and running before showing a status of CREATE_COMPLETE
D. Use the CreationPolicy to ensure it is associated with the EC2 Instance resource

Answer: B and D

Explanation: You might specify additional measures to set up the instance, like installing software packages or Bootstrap applications, for provisioning an Amazon EC2 instance in an AWS CloudFormation stack. Normally, after creating the instance successfully, CloudFormation proceeds with stack creation. However, you can use a Creation Policy so that only after your configuration actions are done, CloudFormation continues with stack creation. You will, therefore, know that your apps are ready to go when the stack is successful.

232. You are a DevOps Engineer in a multi-national company. In order to start building its resources in AWS, the company wants to use CloudFormation templates. The templates for different departments, such as networking, security, apps, etc. are required. What is the best way to develop these templates for CloudFormation?

A. Create separate logical templates, for example, a separate template for networking, security, application, etc. Then nest the relevant templates

B. Consider using Elastic Beanstalk to create your environments since CloudFormation is not built for such customization
C. Consider using OpsWorks to create your environments since CloudFormation is not built for such customization
D. Use a single CloudFormation template, since this would reduce the maintenance overhead on the templates itself

Answer: A

Explanation: When infrastructure grows, common patterns may arise, in which each of your templates declares the same components. You can create

dedicated templates and separate common components. In this way, you can mix and match various templates, but use nested stacks to create a single stack. To create other stacks with in a stack, one should use nested stacks. AWS::CloudFormation::Stack resource is used in the template as reference for other templates to create nested stacks.

233. Mike has an application to support thousands of users, therefore, he has created a DynamoDB table. Every user can access his or her own information only at a particular table is his requirement. Most of the user's accounts are with a third- party ID provider, including Facebook, Google or Amazon Login. How would he implement this? (Choose 2)

A. By using a third-party identity provider such as Google, Facebook or Amazon so users can become an AWS IAM User with access to the application

B. By creating an IAM User for all users so that they can access the application
C. By using web identity federation and register your application with a third-party identity provider such as Google, Amazon, or Facebook
D. By creating an IAM role, which has specific access to the DynamoDB table

Answer: C and D

Explanation: You do not have to create a customized login code or manage your own user identities with a web identity federation. Instead, application users can use well- knownidentity provider IDP for sign-in, such as Amazon Login, Facebook, Google or other OIDC compatible IDPs, and then exchange the received authentication token for AWS temporary security credentials, that map in an IAM role for the resources access in your AWS account. Using an IDP helps you maintain your AWS account secure because your application needs not to integrate and distribute long-term security credentials.

234. Kelvin is designing a CloudFormation stack to create a web server and a database server. He needs to make sure that the database

server is created before the creation of the webserver database. How can he do that?

A. By ensuring that the database server is defined as a child of the web server in the CloudFormation template

B. By ensuring that the database server is defined first and before the web server in the CloudFormation template. The stack creation normally goes in order to create the resources
C. By using the DependsOn attribute to ensure that the database server is created before the web server
D. By ensuring that the web server is defined as a child of the database server in the CloudFormation template

Answer: C

Explanation: You may specify with the DependsOn attribute that a particular resource creation follows another. If you want to restrict the creation of any resource that it should create after the specific resource creation, then add that resource in the DependsOn attribute.

235. You are responsible for designing a number of CloudFormation templates. You must change the stack resources sometimes on the basis of the requirement. How can you monitor the impact of resource change in a CloudFormation stack before stack changes are implemented?

A. By using CloudFormation change sets to check for the impact of the changes

B. By using CloudFormation Stack Policies to check for the impact of the changes
C. There is no way to control this. You need to check for the impact beforehand
D. By using CloudFormation Rolling Updates to check for the impact of the changes

Answer: A

Explanation: When you need to update a stack, it helps you to update stacks

in confidence if you understand how your changes affect running resources before implementing them. Change sets allow you to preview how proposed stack changes can affect your running sources e.g., whether your changes will delete or replace critical resources, and AWS CloudFormation will only make the changes to your stack when you decide to run the change set, allowing you to decide whether to follow the proposed changes or find other changes by creating another change set. The AWS CloudFormation Console, AWS CLI or AWS CloudFormation API allows you to create and manage change sets.

236. Choose 3 answers from the following options, which are true about the OpsWork Stack Instances.

A. You can use instances running on your own hardware

B. You can start and stop instances manually
C. A stacks instance can be a combination of both Linux and Windows based operating systems
D. You can use EC2 Instances that were created outside the boundary of OpsWork

Answer: A, B, and D
Explanation: Following are the features of OpsWork:

- You can start, stop, or automatically scale the number of instances by AWS OpsWorks Stacks. With any stack you can use automatic time-based scaling; Linux stacks can also use load-based scaling
- You can also register instances with a Linux stack that has been created outside of AWS OpsWorks Stacks in addition to the use of AWS OpsWorks Stacks in Amazon EC2 instances. This includes EC2 instances and instances on your own hardware. They have to run one of the Linux distributions that is supported however, you may not be able to register on-premises Windows instances or Amazon EC2
- A stack can run Linux or Windows instances. A stack may have various Linux versions or distributions on various instances, but

Linux and Windows cannot be mixed

237. In order to implement the A/B test strategy for a web application, a DevOps team is using Lambda@Edge. Logic may be added in the CloudFront CDN by Lambda@Edge to select what content to deliver without touching the application code. It is easily possible to test and analyze two versions of app content.

The end user sends the viewer request to CloudFront then it sends the origin request to the origin server. Now the server origin sends the response to cache of CloudFront and then sends the viewer response to end user. From the following, which of the following can be configured to add the logic to choose different versions of applications for A/B testing? (Choose 3)

A. CloudFront Cache

B. Viewer Response
C. Origin Response
D. Viewer Request
E. Origin Request

Answer: C, D, and E

Explanation: In order to choose a version of the application, you can configure Viewer request: When CloudFront receives a viewer request, it checks to see if the object in the edge cache is the requested object.

Origin Request: If the object is in the edge cache, then this function does not execute. It is executed when CloudFront forwards the request to the origin server.

Origin Response: It will execute when the CloudFront receives a response from the origin before caching it to the edge cache. When there is an error from the origin, this function still executes. It will not execute in a case when an object is in the cache or response is generated by the function, which was triggered by an origin request event.

238. A company creates a CloudFormation Template that passes user data to the underlying EC2 Instance. Choose the function that is normally used in the CloudFormation template to transfer data into

the UserData section.

A. "UserData": { "Fn::Ref": {

B. "UserData": { "Fn::GetAtt": {
C. "UserData": { "Fn::FindInMap": {
D. "UserData": { "Fn::Base64": {

Answer: D

Explanation: The intrinsic Fn::Base64 function returns the input string in Base64 representation. This function is typically used to transfer encoded data via the UserData property to Amazon EC2 instances.

239. You are working in a company where CloudFormation stack resources are creating some problems. Select the options from the following, which will help you in debugging. (Choose)

A. Use the AWS CloudFormation console to view the status of your stack

B. Use AWSConfig to debug all the API call's sent by the CloudFormation stack
C. Use CloudTrail to debug all the API call's sent by the CloudFormation stack
D. See the logs in the /var/log directory for Linux instances

Answer: A and D

Explanation: You can view a list of stack events in the AWS CloudFormation stack while your stack is being created, updated, or deleted via console. Select the failure event from this list and then view the status reason for the event. You can see the cloud-init and cfn logs for Amazon EC2 problems. These logs can be found in the /var / log/ directory on the Amazon EC2 instance. These logs record processes and command outputs during instance setup by AWS CloudFormation. View EC2Configure service and cFN logs for Windows in the%ProgramFiles%\Amazon\EC2ConfigService and C:\cfn\log.

240. Kennard has launched multiple instances in different availability zones in an Auto-scaling group as he is running a high traffic

application. He noticed that one availability zone is not receiving any traffic. What could be the reason?

A. Auto-scaling can be enabled for multi AZ only in North Virginia region

B. An availability zone is not added to Elastic Load Balancer
C. Auto-scaling only works in a single region
D. Instances need to be manually added to the availability zones

Answer: B

Explanation: The Elastic Load Balancing creates a load balance node in the available zone when you add an availability zone to your load balancer. Load balancer nodes accept customer traffic and forward requests in one or more availability zones to healthy registered instances.

241. An organization hired a DevOps engineer who is responsible for an AWS Elastic Beanstalk application. They want to move a continuous deployment model, releasing updates to the application multiple times per day with zero downtime. How will the DevOps engineer do this with immediate roll back to the previous version?

A. By creating a second Elastic Beanstalk environment with the new application version, and configuring the old environment to redirect clients using the HTTP 301 response code to the new environment

B. By developing the application to poll for a new application version in code repository; then downloading and installing it to each running Elastic Beanstalk instance
C. By creating a second Elastic Beanstalk environment running the new application version, and swapping the environment CNAMEs
D. By enabling rolling updates in the Elastic Beanstalk environment and setting an appropriate pause time for application startup

Answer: C

Explanation: Due to the fact that Elastic Beanstalk performs an in-place update to your application versions, your application may not be available to users for a short period of time. This down time may be avoided by using a blue/green deployment, in which the new version is deployed to a separate environment, and the CNAMEs of both environments can be switched on instantly to the new version to redirect the traffic.

242. Your company makes you responsible for the development of a number of cloud templates. During a stack update, you must be careful that no one can update production-based resources accidentally on the stack. How can this be done most effectively?

A. By using S3 bucket policies to protect the resources

B. By using MFA to protect the resources
C. By using a Stack based policy to protect the production based resources
D. By creating tags for the resources and then creating IAM policies to protect the resources

Answer: C

Explanation: All update actions are permitted on all resources when a stack is created. Updating all of the resources on the stack can be used (by default) by the one who has stack update permission. During an update, some resources may need to be interrupted or replaced completely, leading to new physical identities or entirely new storage. You can avoid unintended updates or removal of stack resources during a stack update via stack policy. A stack policy is a JSON document defining update actions on designated resources.

243. A successful web product has been built in a fintech company using Node.js with an Elastic Load Balancer instance on EC2 c4.xlarge. And the RDS PostgreSQL m5.4xlarge instance is used for its database. In the AWS region ap-southeast-1, all the AWS services are built. The manager of this company needs to learn how to build a disaster recovery system because this process had not yet been setup. In order to determine the right DR system services, which questions are to be asked first by the manager? (Choose 2)

A. Which type of Elastic Load Balancer does the product use?

B. Whether another engine type of RDS is required?
C. What is the budget for the disaster recovery system?
D. Which instance type of EC2 does it need?
E. What are the RTO and RPO for the product?

Answer: C and E

Explanation: First thing you need to know is the budget because the budget affects the design of the DR system. To achieve the budget target, AWS has offered a number of cost-effective services. The second thing that you need to know is the RTO and RPO. While other options regarding instance type, RDS engine type, and Load balancer type are not important for DR planning.

244. An AWS reader app has been developed by a firm. As the number of users increases, a search feature for the application is increasingly necessary. The firm's team needs this feature to be developed in an independent AWS CodePipeline pipeline as soon as possible. The source stage for the new pipeline is a CodeCommit repository containing the data in the JSON format used for the search service. The build phase has been setup by AWS CodeBuild. Which search feature is best implemented in AWS?

A. Use AWS Lambda to manage an EC2 Elasticsearch application for the search feature. The build stage is responsible for building the Lambda function

B. Use AWS CloudSearch to implement this feature. The build stage uses AWS CLI to configure the data in the search domain of CloudSearch
C. Use AWS EC2 to provision and manage a cluster of servers running SOLR applications for the search feature. The build stage is responsible for baking an AMI
D. Use AWS CloudSearch to implement this search function. The build stage builds a CloudFormation template. Add another deploy stage to deploy the CloudFormation stack for the CloudSearch search domain

Answer: B

Explanation: As this points to the need for a solution as soon as possible, AWS CloudSearch is first to be seen as a fully managed AWS service and users are free to easily set up and manage a website or application search solution. The AWS CLI aws cloudsearchdomain upload-documents can be used to upload JSON format data.

245. Paul has recently developed a new mobile app to handle large-scale analytics workloads stored in Amazon Redshift. Access to the Amazon Redshift tables, therefore, is needed. In practical terms and in security terms, which of the following methods would best allow the tables to be accessed?

A. Create an IAM user and generate encryption keys for that user. Create a policy for RedShift read-only access. Embed the keys in the application

B. Create a RedShift read-only access policy in IAM and embed those credentials in the application
C. Create an HSM client certificate in Redshift and authenticate using this certificate
D. Use roles that allow a web identity federated user to assume a role that allows access to the RedShift table by providing temporary credentials

Answer: D

Explanation: The ideal approach for accessing any AWS service is to use roles. Options A and C are therefore incorrect. You must also use the web identity federation for any web application. Option B is the correct option.

You should request for AWS services, which must be signed with an AWS access key while writing such an app. However, it is recommended that you do not use embedded or long term AWS credentials with apps even in an encrypted store that a user downloads. Instead, build an app that requiresd the AWS temporary security credentials dynamically whenever web identity federation is used The temporary credentials map to an AWS role and only allow executing the tasks required by the mobile app.

246. After the launching of EC2 instance in an Auto-scaling group,

you have implemented a system to dynamically automate configuration deployment and application. Your system uses a configuration management tool, in which the master node is not available. This tool works in a standalone configuration. Because of the application load volatility, new instances should be launched within 3 minutes of the launching of instance operating system. The following times are required for the completion of deployment stages:

- Configuration management agent installation: 2 minutes
- The configuration of an instance using artefacts: 4 minutes
- Application framework installation: 15 minutes
- Deployment of application code: 1 minute

How do you automate deployment with this type of standalone agent configuration?

A. Build a custom Amazon Machine Image that includes the configuration management agent and application framework pre-installed. Configure your Auto-scaling launch configuration with an Amazon EC2 UserData script to pull configuration artefacts and application code from an Amazon S3 bucket, and then execute the agent to configure the system

B. Configure your Auto-scaling launch configuration with an Amazon EC2 UserData script to install the agent, pull configuration artifacts and application code from an Amazon S3 bucket, and then execute the agent to configure the infrastructure and application

C. Build a custom Amazon Machine Image that includes all components pre-installed, including an agent, configuration artefacts, application frameworks, and code. Create a startup script that executes the agent to configure the system on startup

D. Create a web service that polls the Amazon EC2 API to check for new instances that are launched in an Auto-scaling group. When it recognizes a new instance, execute a remote script via SSH to install the agent, SCP the configuration artefacts and application code, and finally execute the agent to configure the system

Answer: C

Explanation: Since new instances must be installed in 3 minutes, all components should be pre-baked in an AMI as a result. When you attempt to use the User Data option, it takes time to install and configure various components based on the time mentioned in the question.

247. An enterprise maintained a large number of applications in AWS via CloudFormation stacks. There are many development teams with various roles such as Developer, UI, QA, etc. To create, modify and delete those stacks, different roles should have different access. In handling all these CloudFormations stacks as a product, the DevOps team needs a centralized AWS resource. And by allowing access to IAM users and groups, the team can handle the product provisioning. Which AWS tool should the team use?

A. AWS Service Catalog

B. Trusted Advisor
C. CloudFormation
D. Systems Manager

Answer: A

Explanation: In AWS Service Catalog, there are some concepts like portfolio and product. A portfolio is the collection of products with configuration detail. The product is a service, which is specified by the AWS CloudFormation template to be made available for deployment on AWS for the user demand.

248. A company has started to use AWS Service Catalog to manage its CloudFormation stacks within the company. In the Service Catalog, several portfolios have been created with relevant products configured. Membership for these various products should be distributed to different users and teams. For example, only the DevOps team may build, change, or uninstall the item for an essential payment service. How should the user access be managed in the Service Catalog?

A. In Service Catalog, use AWS CLI or console to configure ACL

policies to manage user access

B. Assign AWS Cognito User Pools to manage access to different products
C. Assign permissions to IAM users, groups, and roles
D. Make sure the AWS Organization is created. Manage Organization Service Control Policies in Service Catalog to control the access

Answer: C

Explanation: The Service Catalog allows users to access a portfolio that allows them to search the portfolio or launch products within the portfolio. With IAM permissions you can manage access and assign IAM users, IAM groups and IAM roles to the portfolio.

249. Without redirecting or modifying browser URL, the group wants to test various versions of the page for users. The group does not intend to change any front-end or back-end code with additional logic in order to perform this kind of A/B testing. And the CDN network must be made aware of this logic. What is the best place for the switching logic?

A. Lambda and API Gateway

B. Weighted Routing at Route53
C. Global Accelerator at CloudFront
D. Lambda@Edge

Answer: D

Explanation: Without any need to change or redirect the browser URL, Lambda@Edge can provide help to perform A/B testing. So it should be chosen as additional logic to perform A/B testing.

250. Philip is creating templates using CloudFormation that has a parameter, which intakes the database password. How can he make sure that if anybody describes the stack, they will not get the password?

A. By using the hidden property for the parameter value

B. By setting the hidden attribute for the CloudFormation resource
C. By using the NoEcho property for the parameter value
D. By using the password attribute for the resource

Answer: C

Explanation: Set the NoEcho property to true for sensitive parameter values (e.g., passwords). In this way, the value of your parameter is displayed as asterisks (****), whenever anybody describes your stack.

251. An enterprise is using a warm standby DR strategy for its application with RTO and RPO of 30 minutes. The application consists of EC2 instance and RDS MySQL. From the following options, which is helpful to meet these RTO and RPO? (Choose 2)

A. AWS Storage Gateway

B. Route 53
C. RDS MySQL Serverless Feature
D. Auto-scaling Group
E. S3 Glacier

Answer: B and D

Explanation: Route53 is an easy, fast and safe solution to help in switching the traffic to standby when needed. The auto-scaling group is appropriate to meet the required capacity by adjusting the number of instances.

252. A big company handles the source code using AWS CodeCommit. There are over 100 repositories and numerous teams working on various projects. You need to provide the appropriate access to different users as a DevOps engineer. For example, the development team should not access the repositories containing sensitive information. How are you supposed to manage this?

A. In IAM, for each IAM user, add an IAM policy that allows or denies actions on repositories based on repository names

B. Configure Git tags in AWS CodeCommit repositories. Create policies in IAM that allow or deny actions on repositories based on the Git tags

C. Tag repositories in AWS CodeCommit. Create policies in IAM that allow or deny actions on repositories based on the tags associated with repositories
D. In AWS CodeCommit console, create CodeCommit policies to IAM groups that allow or deny actions on repositories

Answer: C

Explanation: In AWS CodeCommit, you have a new feature to tag the repositories because by using tags, IAM policies can easily be configured for controlled access. For example, you define the policy of denying permission to all the resources to perform CodeCommit on the sensitive data.

253. John's company has started using AWS CodeCommit to maintain the source code. Thousands of developers are present and new Git repositories are introduced regularly in CodeCommit. The DevOps team was instructed to formulate an approach in the creation of a new repository through the running of a script and after the script is running, it can automatically add setup code, such as model and README. What is the best approach to manage everything?

A. Put the code into the S3 bucket. Design a CloudFormation template to create a repository with the resource type AWS::CodeCommit::Repository. Add a property to include the code. Create a CloudFormation stack when needed

B. CodeCommit does not support this. Use GitHub instead
C. Design a CloudFormation template to create a repository with the resource type AWS::CodeCommit::Repository. Design a Lambda function to commit the initial code to the new repository
D. Design a shell script using Git commands to create a new repository and submit the init code

Answer: A

Explanation: When you are creating a repository with AWS CloudFormation, AWS CodeCommit supports it including the code. CloudFormation stack can be handled easily. A new repository is configured via stack and in the meantime, the program code is inserted in the repository.

254. You are using AWS Elastic Beanstalk to run a social media marketing application, for which you have written a component in Ruby. To support different marketing campaigns, this application component sends messages to social media sites. Your management needs to record responses to these messages from social media to analyze the marketing campaign's effectiveness compared to past and future efforts. A new application component has already been developed for the social media site API to read the replies. Which method should you use to record social media responses for analytical historical data in a sustainable data storage that can be accessed at any time?

A. Deploy the new application component as an Elastic Beanstalk application, read the data from the social media sites, store it in DynamoDB, and use Apache Hive with Amazon Elastic MapReduce for analytics

B. Deploy the new application component in an Auto-scaling group of Amazon EC2 instances, read the data from the social media sites, store it with Amazon Elastic Block Store, and use AWS Data Pipeline to publish it to Amazon Kinesis for analytics

C. Deploy the new application component in an Auto-scaling group of Amazon EC2 instances, read the data from the social media sites, store it in Amazon Glacier, and use AWS Data Pipeline to publish it to Amazon RedShift for analytics

D. Deploy the new application component as an Amazon Elastic Beanstalk application, read the data from the social media site, store it with Amazon Elastic Block store, and use Amazon Kinesis to stream the data to Amazon CloudWatch for analytics

Answer: A

Explanation: For all applications requiring a consistent single-digit millisecond latency at every scale, then Amazon DynamoDB is the best data service. Amazon DynamoDB is a fast and flexible NoSQL database service. It is a cloud database, which is fully managed and supports key value storage and document models. Due to its flexible data model, reliable performance and automated throughput capacity scaling, the system fits mobile, web,

gaming, ad technology, IoT and many other applications greatly.

255. A start-up IT firm has just moved its most important web-based software to AWS. A disaster recovery program in AWS must be built as soon as possible, otherwise, the failure may have a huge impact on the credibility of the business and on the cash flow. The organization does, however, have a budget run out and has to review its operating costs. What tools will help the organization control costs during the implementation of a disaster recovery system? (Choose 2)

A. Use an S3 lifecycle policy to move all stored data in S3 to Glacier after one day to lower down the cost

B. Use a Trusted Advisor to monitor the EC2 instances that have a low utilization rate in the standby system. Terminate those instances to save costs
C. Use a suitable Auto-scaling Group to control the number of running instances
D. Configure an EBS Lifecycle Policy to delete old EBS snapshots
E. Create more spot instances in the hot standby system

Answer: C and D

Explanation: In order to reduce costs, ensuring that the process can operate as usual and meet the required RPO and RTO must also be taken into account.

The Auto-scaling group is helping to reduce the number of operating instances in order to save costs. EBS Lifecycle Policy promotes effective handling of EBS snapshots. Cost-saving can be done by removing old images.

256. You maintain an application that serves as a front end and uses MongoDB for document management on an acceptable web server. You use the user data tab to set up the application to pre-bake AMI with the newest version of the web server. You now have an amendment to the underlying version of the OS and must, therefore, deploy it. How will you do that as easily as possible?

A. By creating a CloudFormation stack with the new AMI and then deploying the application accordingly

B. By creating a new pre-baked AMI with the new OS and using the User Data section to deploy the application
C. By creating a CloudFormation stack with the new AMI and then deploying the application accordingly
D. By creating an Opswork stack with the new AMI and then deploying the application accordingly

Answer: B

Explanation: The ideal way is to continue the same deployment process that was used previously and create a new AMI and use the user data section in order to deploy the application.

257. Your development team uses Elastic Beanstalk. Deploying multiple versions of your application is your responsibility. How do you ideally ensure that you do not cross Elastic Beanstalk's application version limit?

A. By using lifecycle policies in Elastic Beanstalk
B. By creating a script to delete the older versions

C. By creating a Lambda function to delete the older versions
D. By using AWSConfig to delete the older versions

Answer: A

Explanation: Elastic Beanstalk creates an application version whenever you upload your application's new version with the Elastic Beanstalk console or EB CLI. You will eventually reach the application version limit and will not be able to create a new version of that application if you do not delete previously created versions that you are no longer using. The application version lifecycle policy can prevent you from reaching the limit. An Elastic Beanstalk lifecycle policy tells Elastic Beanstalk to remove old versions of an app or remove versions of the application if the total number of versions of the app exceeds the specified number.

258. You are working on an application that has three EC2 instances in the Auto-scaling group behind an Elastic Load Balancer. You find

that the Auto-scaling group was updated with a new launch configuration that refers to an updated AMI. Your ELB health checks were successful yet, you still have received complains of error from users. What will you do to prevent this situation from happening again?

A. Update the launch configuration instead of updating the Auto-scaling group

B. Create a new ELB and attach the Auto-scaling group to the ELB
C. Manually terminate the instances with the older launch configuration
D. Create a new launch configuration with the updated AMI and associate it with the Auto-scaling group. Increase the size of the group to six and when instances become healthy, revert to three

Answer: D

Explanation: At a time, one launch configuration is associated with an Auto-scaling group, and after creation, you cannot change the launch configuration. You can use an existing launch configuration as the foundation for a new launch configuration then, upgrade the Auto-scaling group to use a new launch configuration. After changing the launch setup for an Auto-scaling group, any new instances can be launched with the new configuration options without affecting the existing instances. Then, to check the launching of new instances, change Auto-scaling group size to 6 and once the instances are initiated, turn it back to 3.

259. Choose the invalid statement regarding the application deployment if you want to deploy applications to ELB.

A. The application can be bundled in a zip file

B. The application should not exceed 512 MB in size
C. The application can include parent directories
D. The application can be a war file, which can be deployed to the application server

Answer: C

Explanation: If a new application or application version is to be deployed using the AWS Elastic Beanstalk console, a source package is needed. The following requirements must be met by your source bundle:

- It should consist of one ZIP or WAR file (multiple WAR files can be included within your ZIP file)
- It should not exceed 512 MB
- It should not include a parent folder or top-level directory (subdirectories are fine)

260. In AWS CodeDeploy, ABC company used a single web application, which includes EC2, for AWS services. CodeDeploy works well, but your manager wants to get notifications in time when the next installation is successful or fails, in the slack channel. You were told to implement this in AWS CloudWatch events. A Lambda feature that only covers deployments would be the target of the Event rule. What is the right event pattern for the rule?

A. {

```
"source": [
    "aws.codedeploy"
],
"detail-type": [
     "CodeDeploy Deployment State-change Notification"
],
"detail": {
     "state": [
          "FAILURE",
           "SUCCESS"
          “READY”
          “START”
             ]
      }
}
```

B. {

"source": [
"aws.codedeploy"
],
"detail-type": [
"CodeDeploy Deployment State-change Notification"
],
"detail": {
"state": [
"FAILURE",
"SUCCESS"
]
}
}

C. {

"source": [
"aws.codedeploy"
],
"detail-type": [
"AWS API Call via CloudTrail”
],
"detail": {
"state": [
"FAILURE",
"SUCCESS"
]
}
}

D. {

"source": [
"aws.codedeploy"
],
"detail-type": [

```
        "CodeDeploy Deployment State-change Notification"
    ],
    "detail": {
        "state": [
            "FAILURE",
             "START"
               ]
         }
    }
```

Answer: B

Explanation: The CloudWatch Events policy is useful for tracking CodeDeploy status. Because this example needs to track the change in deployment status, the “detail-type” should be "CodeDeploy Deployment State-change Notification".

261. Which of the following 3 things would you be able to accomplish with the benefit of the CloudWatch logs?

A. Send the log data to AWS Lambda for custom processing or to load into other systems
B. Record API calls for your AWS account and deliver log files containing API calls to your Amazon S3 bucket
C. Stream the log data to Amazon Kinesis
D. Stream the log data into Amazon Elasticsearch in near real-time with CloudWatch logs subscriptions

Answer: A, C, and D

Explanation: For fast and continuous data intake and aggregation, Amazon Kinesis can be used. The data used include IT infrastructure log data, application logs, social media, feeds for market data, and clickstream data.

Amazon Lambda is a web service that can be used to compute logs published by CloudWatch logs without servers.

To deploy, operate, and scale Elasticsearch for log analytics, full text search, application monitoring, and many more in the simplest way, Amazon Elasticsearch Service can be used.

262. In your organization, you have set up the following AWS

services: Auto-scaling Group, Elastic Load Balancer, and EC2 Instances. If the utilization of CPU is less than 30%, you have to terminate an instance from the Auto-scaling group. How will you do that?

A. By creating a CloudWatch alarm to send a notification to the Auto-scaling group when the aggregated CPU utilization is less than 30% and configuring the Auto-scaling policy to remove one instance

B. By creating a CloudWatch alarm to send a notification to SQS. SQS can then remove one instance from the Auto-scaling Group
C. By creating a CloudWatch alarm to send a notification to the admin team. The admin team can then manually terminate an instance from the Auto-scaling group
D. By creating a CloudWatch alarm to send a notification to the ELB. The ELB can then remove one instance from the Auto-scaling Group

Answer: A
Explanation: You should define two policies, one for scaling in (terminating instances) and one for scaling out (launching instances) for monitoring each event. For example, when the network bandwidth reaches a certain level, you want to scale out and for that purpose, you have to create a policy specifying that Auto-scaling should start with a certain number of instances to help your traffic. But, if the network bandwidth level goes down when the network bandwidth level goes back down, you have to define the scale in the policy.

263. What is not a supported Elastic Beanstalk service platform?

A. PHP
B. AngularJS
C. .Net
D. Java

Answer: B
Explanation: Following are the supported platforms on Elastic Beanstalk:

- Go
- Java SE
- Java with Tomcat
- .NET on windows server with IIS
- Node.js
- PHP
- Python
- Ruby
- Packer Builder
- Single container Docker
- Multicontainer Docker
- Preconfigured Docker

264. Your team will run a Java project using AWS CodeStar. AWS CodeCommit includes the source code. The build stage is managed with AWS CodeBuild followed by a stack with resources like the Lambda function via the CloudFormation deployment stage. Several team members, including Jason (owner), Tony (viewer) and Eric (contributor), have been added to the CodeStar project. The various roles in the project should be allowed to different members. How are CodeStar's permissions managed?

A. In CodeStar console, configure the team member roles by assigning different read and write permissions to stages such as build or deploy

B. For different team members, users need to create appropriate IAM policies first and then assign these policies to IAM users
C. Different team member roles have relevant IAM policies allocated automatically by AWS. CodeStar users only need to make sure the correct roles are assigned to team members
D. Users need to create IAM service roles with suitable IAM policies. Then assign these service roles to different team

members in CodeStar depending on their roles

Answer: C

Explanation: When a project is created, AWS CodeStar will automatically create IAM policies for your customers. These rules are used to control levels of access to CodeStar team members. The related IAM policies are automatically added to the IAM users when users add team members with roles.

265. In a project, a Linux-based instance stack in Opswork has been defined. Furthermore, you want to attach a database to it. Which of the following is an important step towards ensuring that the Linux application can communicate with the database?

A. Configuring SSL so that the instance can communicate with the database
B. Configuring database tags for the OpsWork application layer
C. Adding another stack with the database layer and attaching it to the application stack
D. Adding the appropriate driver packages to ensure the application can work with the database

Answer: D

Explanation: For Linux Stacks, you have to add the appropriate driver package to the associated application layer if you want to associate an Amazon RDS service layer with your application. To do this, follow the given steps:

i. Click “Layers” in the navigation pane and open the app server's “Recipes” tab.
ii. Click “Edit” and insert OS Packages with the appropriate driver package. For example, if a layer contains instances of Amazon Linux or mysql-client, you should specify mysql if the layer contains instances of Ubuntu.
iii. Save changes and redeploy the app.

266. You have been tasked to build a dashboard for security control

in AWS. The dashboard should be able to determine if EC2 instances are vulnerable and exposed (CVEs). This incident should be detected, for example, when an EC2 instance does not download a particular patch and is exposed to a known CVE. What is the best approach to do so?

A. Configure AWS Macie and include CVE rule package in the assessment template

B. Enable AWS Inspector and make sure all EC2 instances have the Inspector agents installed properly. Include the CVE rule package in the assessment template
C. In AWS Systems Manager, include CVE patches in patch baselines. Use patch manager to apply system patches to all EC2 instances
D. Enable AWS GuardDuty and include the CVE rule package in the GuardDuty template. Monitor CVE findings in the console

Answer: B

Explanation: One of the rule packages that AWS Inspector can configure is common faults and exposures (CVEs). & nbsp; AWS Inspector can detect if EC2 instances are exposed to CVE after configuring the Common Vulnerabilities and Exposures (CVE) rule packages on the evaluation template.

267. Your company designed an application to transfer all user logs to a Kinesis channel, and those logs will live for 24 hours in the streams. A recent strategy of Splunk Enterprise is to scan, track, and review request logs. Also, on an EC2 instance, the Splunk Server has been deployed. What is the best approach to loading streaming data into the Splunk instance on the Kinesis Stream?

A. Use AWS SDK in an EC2 instance to get the records from Kinesis Stream and forward the records to the Splunk instance

B. Use Amazon Kinesis analytics to analyze and transfer real-time streaming data in Kinesis Stream to the destination, which is the Splunk instance

C. Use Amazon Kinesis Data Firehose as a fully managed service to deliver real-time streaming data in Kinesis Stream to the Splunk instance
D. Configure the Kinesis Stream to auto deliver the received logs to the Splunk destination for the server to index the logs

Answer: C

Explanation: Amazon Kinesis Data Firehose is an appropriate data consumer for the data supplier, in this case, the Kinesis Stream. Kinesis Firehose is used to send logs to destinations such as Amazon S3, Amazon Redshift, Splunk, and Amazon ES.

268. You are creating a DynamoDB application for the storage of JSON data. The read and write capability of the DynamoDB table has already been set. The amount of traffic received during the deployment period by the application is not known. Your IT officer asks you to ensure that DynamoDB is not throttled and is not an application bottleneck. What steps should you take for this? (Choose 2)

A. Monitor the SystemErrors metric using CloudWatch

B. Create a CloudWatch alarm, which would then send a trigger to AWS Lambda to increase the Read and Write capacity of the DynamoDB table
C. Create a CloudWatch alarm, which would then send a trigger to AWS Lambda to create a new DynamoDB table
D. Monitor the ConsumedReadCapacityUnits and ConsumedWriteCapacityUnits metric using CloudWatch

Answer: B and D

Explanation: ConsumedReadCapacityUnits and ConsumedWriteCapacityUnits over the specified time period can monitor for a DynamoDB table to track the usage of your provisioned throughput.

269. You are responsible for an application, which is using EC2, ELB, and Auto-scaling. The ELB access logs were requested by your manager. You do not find any log in the S3 bucket. Why is that

happening?

A. The Auto-scaling service is not sending the required logs to ELB

B. You do not have the necessary access to the logs generated by ELB
C. The EC2 Instances are not sending the required logs to ELB
D. By default, ELB access logs are disabled

Answer: D

Explanation: Access logging is by default disabled. It is an Elastic Load Balancing feature. Once the load balancer logs are enabled, Elastic Load Balancing captures the logs and keeps them in a specified Amazon S3 bucket. The access logging can be disabled at any time.

270. For several years, an organization has been using the Jenkins database on site as a CI/CD tool. Most of its services have recently moved to AWS. AWS CodePipeline will be used as the new tool in order to replace the Jenkins database by the DevOps team. Initial research and reporting were given to the CEO. What are CodePipeline's strengths relative to that of the Jenkins server? (Choose 2)

A. IAM policies can be configured to control the access to CodePipeline resources

B. Similar to Jenkins, CodePipeline also provides a large number of plugins in the AWS marketplace
C. AWS CodePipeline is totally free and you only need to pay related resources generated in the pipeline such as EC2
D. AWS CodePipeline integrates well with other AWS services such as CodeCommit and CodeBuild. It is easily configured for users of the AWS ecosystem
E. AWS CodePipeline is open-sourced just as Jenkins so that users can contribute to the community

Answer: A and D

Explanation: You can use AWS CodePipeline with various other services because it will integrate easily with them and also be easily configured. In order to manage the resource access in the CodePipeline, IAM policies can be used for users, groups, and roles.

271. Which of the following can be used to monitor whether the changes made to your AWS resources are reliable and durable logging solutions?

A. Create a new CloudTrail with one new S3 bucket to store the logs. Configure SNS to send log file delivery notifications to your management system. Use IAM roles and S3 bucket policies on the S3 bucket that stores your logs

B. Create three new CloudTrail trails with three new S3 buckets to store the logs one for the AWS Management console, one for AWS SDKs and one for command line tools. Use IAM roles and S3 bucket policies on the S3 buckets that store your logs
C. Create a new CloudTrail trail with an existing S3 bucket to store the logs and with the global services option selected. Use S3 ACLs, and Multi Factor Authentication (MFA) Delete on the S3 bucket that stores your logs
D. Create a new CloudTrail trail with one new S3 bucket to store the logs and with the global services option selected. Use IAM roles S3 bucket policies and Multi Factor Authentication (MFA) Delete on the S3 bucket that stores your logs

Answer: D

Explanation: AWS CloudTrail service is integrated with AWS Identity and Access Management (IAM). It is a service that logs AWS events from or on behalf of your AWS account. CloudTrail logs AWS authenticated API calls as well as AWS sign-in events and collects information of this event in the files that are supplied to Amazon S3 buckets. You must make sure all services are included. Therefore, option C is partly correct.

Options A and B are incorrect because having three S3 buckets and SNS notifications just adds overhead.

272. A crew is working on the development of a new AWS EC2 web

application. Atlassian JIRA Software is used for many existing projects within the company. The team is expected to work with a single project management portal where information is available such as JIRA tales, on-going implementation, system endpoints, etc. for this new project. How can the team meet the expectations?

A. Create an AWS CodeStar project. Customize the project dashboard as required with a link to the JIRA server URL

B. Create an AWS CodeStar project. Configure the connection to JIRA in the CodeStar project. Customize the project dashboard as required

C. Use AWS CodePipeline to manage the source, build and deployment. Add a stage with the action provider as Jenkins. In the Jenkins server, configure a JIRA plugin to integrate with Atlassian JIRA

D. Use the CloudFormation stack to create CloudWatch dashboards to manage the project including the JIRA URL

Answer: B

Explanation: AWS CodeStar has provided compatibility with the Atlassian JIRA framework, which allows the CodeStar dashboard to easily integrate issue tracking and project management tools that the JIRA provides. After JIRA has been incorporated, all JIRA project tickets can be easily seen by CodeStar users. In the dashboard, you can find JIRA information.

273. A large company has a massive quantity of data in AWS S3. These S3 buckets include the application's read or write data. The safety auditor was concerned that some sensitive information could be revealed in S3. Some programs, for example, can store some text files containing PII data for customers. In these S3 buckets, the auditor asked for a solution to quickly scan potential security issues. What is the best solution?

A. Configure Amazon Athena in S3 and create Athena SQL tables. Query security issues by using SQL commands

B. Enable AWS GuardDuty as it can analyze all the application

data in S3 and generate security findings

C. Configure AWS Inspector in S3. It is able to use machine learning to search for security issues in S3 and provide CloudWatch alarms to the admin users
D. Enable Amazon Macie as it can scan security issues in S3 and generate alerts based on the level of risks

Answer: D

Explanation: When you are concerned about the PII information security in Amazon S3, always think about Macie. After enabling it, S3 objects can be scanned by using PII priority.

274. You have an OpsWork stack on Linux instances. Your recipe execution failed. How can you diagnose the reason for failure?

A. Bu logging into the instance and checking if the recipe was properly configured

B. By de-registering the instance and checking the EC2 Logs
C. By using AWS CloudTrail and checking the Opswork logs to diagnose the error
D. By using AWS Config and checking the Opswork logs to diagnose the error

Answer: A

Explanation: Failure to use a recipe will lead to the instance setup failed state rather than online. While the instance of AWS OpsWorks Stacks is not online, it is often useful to login to resolve the matter in EC2. The EC2 instance is running. You can check whether an application or a personalized cookbook is installed correctly. The AWS OpsWork Stacks built-in support for SSH and RDP login is available only for the online state instances.

275. AWS platform is used by the company to host the bulk of its applications and services. You have handled a great many AWS assets as an AWS administrator to meet business needs. In many AWS regions and accounts, you sometimes need to build, upgrade or delete similar resources. For example, to decide whether CloudTrail is allowed on all accounts, you have to set up an AWS Config rule.

What is the safest way to implement these settings in multiple regions or on account of single operations?

A. Create CloudFormation StackSets using templates. Specify regions and accounts depending on the requirements for the StackSets

B. Configure an AWS CodePipeline to deploy AWS resources in a deployment stage via CodeDeploy. Execute the pipeline in different regions and accounts
C. Create Lambda functions to use the AWS SDK to create AWS resources. Run the Lambda functions in different regions and accounts
D. Prepare CloudFormation templates and create CloudFormation stacks using the templates in various regions and accounts

Answer: A

Explanation: The key question is that AWS infrastructure should be set up or managed in a single operation across different regions and accounts. CloudFormation StackSets must first be considered as CloudFormation stacks can be easily supplied to selected target accounts in specified regions. For a single CloudFormation StackSet, regions, and accounts can be picked.

276. The Kinesis Stream was used by a financial firm to store processed logs from a busy application in real time. The data is then sent to a Kinesis Firehose distribution system that provides information to the final destination of the S3 container. The data input format is RFC3163 Syslog. Before the data is delivered, the format must be converted to JSON in Kinesis Firehose. How will this be done?

A. Kinesis Firehose cannot transform the data format inside of it. Instead, it has to be done in Kinesis Stream

B. Configure Kinesis Data Firehose to use third-party JSON deserializer tool Apache Hive JSON SerDe to convert the data to JSON format
C. Create a Lambda function for data transformation using a

blueprint. Kinesis Data Firehose can invoke the Lambda function to transform incoming source data

D. In Kinesis Data Firehose, invoke AWS Glue to create a schema in order to modify the format of incoming source data

Answer: C

Explanation: In order to convert the data format in Kinesis Firehose, you can use a blueprint that converts the syslog data into the JSON format.

277. The DevOps group assesses AWS CodeBuild to figure out if it is suitable for developing new software. Different objects or environments are needed to build artifacts for these applications. The environmental image, for which the construction project is to run must be selected in AWS CodeBuild. From the given options, what are good CodeBuild environment images? (Choose 2)

A. A Centos Linux image from an AMI created by the user

B. A custom docker image, which is hosted in an external Docker registry

C. A Suse Linux image managed by the user

D. A RedHat Linux image managed by AWS CodeBuild

E. A Windows server image managed by AWS CodeBuild

Answer: B and E

Explanation: For CodeBuild build environments, there are two types of environments: custom or managed images. Instead of AMI, Docker image should be used. RedHat Linux image is not supported and custom image should be a docker image. While Windows server image is supported.

278. Your organization owns several AWS accounts. The AWS operation team creates several base docker images in AWS ECR. Another development team is working on a new project, in which the build phase needs to use AWS CodeBuild to build artifacts. One requirement is that the environment image of CodeBuild must use an ECR docker image owned by the operation team. However, the ECR docker image is located in a different AWS account. How would you resolve this and create the CodeBuild project?

A. Select custom image and choose another registry for an external Docker registry. In external registry URI, enter the ECR repository URI of the other account

B. Select the custom image and choose the ECR image registry. Enter the full ECR repository URI for the repository in the other account
C. CodeBuild does not support cross account ECR images. Copy the image to the ECR registry in the same account first
D. Select AWS managed Ubuntu image. In the image, pull the ECR docker image from another account

Answer: B

Explanation: AWS CodeBuild supports accessing cross-account ECR images. In the console of AWS CodeBuild, you can select the ECR registry and then select the account.

279. Robert's team has created a new AWS CodeDeploy software to simplify Amazon EC2 deployments under auto-scaling. A deployment operation may have failed because of software problems. In the production environment, your manager asked you to automatically setup a rollback to the deployment group. What statement is correct when it comes to the automatic rollback setup of CodeDeploy?

A. Automatic rollback can only be configured for EC2 instances. On-premises instances do not have this feature

B. Users can configure SNS notifications for deployment activities. Rollback can be triggered whenever SNS topics receive notifications
C. Automatic rollback can be triggered when alarm thresholds are met
D. There is no automatic rollback for the CodeDeploy deployment group. However, users can trigger rollback manually

Answer: C

Explanation: By configuring automatic rollback in AWS CodeDeploy, it can

be triggered in case of deployment fails or if a certain threshold is met.

280. A multi-level architecture is being operated on AWS with the Nginx web server instances. The application is having bugs when operated by users. How can you quickly and effectively identify those errors?

A. Send all the errors to AWS Lambda for processing

B. Send all the errors to AWS Config for processing
C. Install the CloudWatch Logs agent and send Nginx access log data to CloudWatch. From there, pipe the log data through to a third party logging and graphing tool
D. Install the CloudWatch Logs agent and send Nginx access log data to CloudWatch. Then, filter the log streams for searching the relevant errors

Answer: D
Explanation: For searching and matching terms, phrases or values in your log events, you can use metric filters. In your log events, you can increase the value of a CloudWatch metric when a metric filter finds a term, phrase or values. For example, to scan and count the occurrence of the word ERROR in your log events, a metric filter can be created.

281. A new project to build a pipeline for a new Android app was assigned to you in the AWS CodePipeline operation. The source stage of the pipeline is GitHub and CodeBuild is used for building the app. This uses a buildspec file to create objects that are stored in an S3 bucket during the construction phase. To test the new version of the App in AWS Device Farm, a further step needs to be added. The QA team has already developed a test project in AWS Server Farm. In AWS CodePipeline, how should you configure this new stage?

A. Add a new test stage and add AWS Device Farm as the action provider. Configure the AWS Device Farm project ID and device pool ARN in the stage

B. Add a new stage and add an SNS topic as the action provider.

The SNS topic will trigger the AWS Device Farm project to execute the test

C. Add a new deploy stage and add AWS CloudFormation as the action provider
D. The CloudFormation stack will create and initiate the AWS Device Farm project
E. Add a new stage and add an action provider using a Lambda function. The Lambda function is responsible to trigger the AWS Device Farm project

Answer: A

Explanation: You can create a continuous integration flow using the AWS CodePipeline, in which your software will be built and tested every time you move a commit into your repository.

Add a new test phase and include the service provider AWS Server Farm. Configure the AWS Device Farm project ID and device pool ARN in the stage.

282. As a DevOps Engineers, you have to host a custom application with custom dependencies for a development team by using AWS service. From the following options, choose the perfect way to perform your task.

A. Package the application and dependencies with Docker, and deploy the Docker container with Elastic Beanstalk

B. Package the application and dependencies with in Elastic Beanstalk, and deploy with Elastic Beanstalk
C. Package the application and dependencies with Docker, and deploy the Docker container with CloudFormation
D. Package the application and dependencies in an S3 file, and deploy the Docker container with Elastic Beanstalk

Answer: A

Explanation: The deployment of a web application from Docker containers is supported by Elastic Beanstalk. You can set your own runtime environment with Docker containers. You can choose a platform, programming language, and any application dependencies that are not

supported by other platforms, such as package managers or tools. Docker containers are autonomous and contain all the configuration information and software needed to run your web application.

283. Raffaele has heavy AWS users and possesses many AWS resources such as EC2, S3, RDS, etc. in his company. Now he is required to develop an ongoing monitoring service in AWS to track services from the safety standpoint. For example, if an EC2 instance is compromised and used to carry out a Denial of Service (DoS) attack using the UDP protocol, the service should be able to identify potential risks. What should Raffaele do?

A. Activate VPC Flow Logs, AWS CloudTrail event logs, and DNS logs and transfer the logs to a dedicated S3 bucket. Configure Athena to query the logs to identify potential security problems

B. Enable AWS Enterprise support plan and activate full features of Trusted Advisor, which can quickly provide alarms for security related issues
C. Enable AWS Macie to continuously scan AWS security risks in resources such as EC2. It can identify potential issues and provide alarms such as if an EC2 instance is compromised
D. Enable AWS Macie to continuously scan AWS security risks in resources such as EC2. It can identify potential issues and provide alarms such as if an EC2 instance is compromised

Answer: D
Explanation: Amazon GuardDuty is a security service capable of constantly tracking and analyzing the information assets including VPC stream logs, AWS CloudTrail event logs, and DNS logs. Like a high-risk alarm found by GuardDuty when an EC2 instance may be used to attack Denial of Service using the UDP Protocol on TCP port.

284. In his AWS EC2 instances, Raymond recently experienced an IT security incident. An offender used an EC2 penetration testing tool from Kali Linux, found weaknesses in the EC2 configurations, and gained unauthorized access to the company's resources. To ensure

that all EC2 cases are properly patched, Raymond needs to develop a plan. For such potential security risks, a monitoring tool is also needed. How should he work together with his team to fulfill the requirements? (Choose 2)

A. Configure monitoring dashboard in AWS QuickSight, which uses machine learning skills to discover security incidents that are happening

B. Use AWS Systems Manager Run Command to apply necessary patches every 30 days to ensure all EC2 instances are always patched compliant
C. Enable AWS GuardDuty to monitor potential security incidents. Create CloudWatch Event rules based on the findings and trigger SNS notifications
D. Configure patch baselines in AWS Systems Manager and use Patch Manager to apply patches in a maintenance window
E. Configure AWS Macie to continuously monitor security issues for AWS resources. Configure SNS notifications based on Macie alarms in CloudWatch Events

Answer: C and D

Explanation: In EC2 instances, patches can easily be applied according to defined patch baselines by the patch manager. Using AWS GuardDuty to track the security problems listed above is the best solution. PenTest: IAMUser / KaliLinux is a finding type of GuardDuty. When a computer operating Kali Linux makes API calls with your AWS Account credentials, GuardDuty will report this threat.

285. Rex has several AWS CloudFormation StackSets. A StackSet has been developed in several regions to set up web application network resources, such as IAM roles and Security Groups. A parameter value of CloudFormation StackSets must be modified because of some new features in the program. Nevertheless, only two regions need this modification and the parameter for the other regions should not be modified. How can he do this?

A. Choose Override StackSet parameters from the Actions menu.

Specify the two regions that he wants to modify and then override the StackSet parameters for these regions

B. Specify the two regions that he wants to change and then modify the StackSet parameters for only these two regions
C. Parameters of StackSets cannot be modified for selected regions. Deregister these two regions from the StackSets and then create a new StackSet with new parameters in these two regions
D. Parameters of StackSets can only be modified for all regions. This operation cannot be done unless the CloudFormation stacks in these two regions are removed from the StackSets

Answer: A

Explanation: In this case, in two regions, stack instances need to have property values different from the one specified when StackSets are created. The parameters for current stack instances can be overridden. First, select the StackSets and choose to override StackSet parameter then specify the account and region. Edit the required parameter value.

286. Your application is running behind a load balancer on Amazon EC2 instances. Your company has decided to use a strategy for the blue/green deployment. How do you do this for every deployment?

A. Launch more Amazon EC2 instances to ensure high availability, de-register each Amazon EC2 instance from the load balancer, upgrade it, and test it, and then register it again with the load balancer

B. Set up Amazon Route53 health checks to fail over from any Amazon EC2 instance that is currently being deployed to
C. Using AWS CloudFormation, create a test stack for validating the code, and then deploy the code to each production Amazon EC2 instance
D. Create a new load balancer with new Amazon EC2 instances, carry out the deployment, and then switch DNS over to the new load balancer using Amazon Route53 after testing

Answer: D
Explanation: To do this you must:

- Firstly, create a new ELB to show new changes in production
- For the distribution of traffic to the 2 ELB based on an 80- 20% traffic scenario, use the Weighted Route Policy for Route53. This is the normal scenario, according to the requirement, the percentage can be changed
- Finally, if all modifications have been tested, Route53 can be set to 100% for the new ELB

Option A is incorrect, as the deployment scenario is not blue green. You are not able to control the users for a new EC2 instance.
Option B is incorrect because this is not a Blue Green Deployment Failure scenario. You need to have two environments working side by side in Blue Green deployments.
Option C is wrong because the changes will run side by side with a production stack.

287. The team is creating an online bid program that has been used to store information for customers in DynamoDB tables. For read, a response time of microseconds is needed. The latest DynamoDB tables do not seem to provide this efficiency however. What is the best approach to greatly improve the reading output without altering current program logic?

A. Configure a DynamoDB Accelerator (DAX) cluster for the application to use which can deliver up to a 10x performance improvement

B. Create a read replica table in another region (Global Table) to improve the read capacity
C. Create a Global Secondary Index for the DynamoDB table so that queries can be more efficient
D. Configure an on-demand read/write capacity mode for the DynamoDB table

Answer: A

Explanation: DAX is a microsecond latency DynamoDB service for access to eventually consistent data. DAX is a fully managed, highly available in-memory cache that boosts DynamoDB performance x10. However, because DAX is compatible with existing DynamoDB API calls, no alteration in application code is required.

288. You decided to change the instance type of your production instances that run in the Auto-scaling group. To launch your architecture, you used the CloudFormation template and currently used 4 instances in production. The service cannot be interrupted therefore two instances should always run during the update. Which of the following options can be applicable?

A. AutoScalingReplacingUpdate
B. AutoScalingScheduledAction
C. AutoScalingRollingUpdate
D. AutoScalingIntegrationUpdate

Answer: C

Explanation: The AWS::AutoScaling::AutoScalingGroup resource supports an UpdatePolicy attribute, which defines how an Auto-scaling group resource is updated when an update to the CloudFormation stack occurs. A common approach is executed rolling update for updating an Auto-scaling group by defining the AutoScalingRollingUpdate policy. This keeps the same Auto-scaling group and, according to the indicated parameters, replaces old instances with new ones.

289. You configure the Continuous Integration (CI) system to create AMIs in your deployment process. You want to build them in a cost-effective, automated way. What method are you supposed to use?

A. Upload all contents of the image to Amazon S3 launch the base instance, download all of the contents from Amazon S3 and create the AMI
B. Have the CI system launch a new spot instance bootstrap the code and apps onto the instance and create an AMI out of it

C. Have the CI system launch a new instance, bootstrap the code

and apps onto the instance and create an AMI out of it

D. Attach an Amazon EBS volume to your CI instance, build the root file system of your image on the volume, and use the CreateImage API call to create an AMI out of this volume

Answer: B

Explanation: You can add Automation as a post-build step to pre-install application releases to Amazon Machine Images (AMI) if you use Jenkins software within a CI/CD pipeline. The Jenkins scheduling function can also be applied to call Automation and create your own OS patching cadence.

290. A website is running in a virtual private cloud, using a load balancer and an Auto-scaling group. Your Head of Security has asked you to set up a system of monitoring that will rapidly detect and notify your team when there is a sudden increase in traffic. How would you configure that?

A. Set up an Amazon CloudWatch alarm for the Elastic Load Balancing NetworkIn metric and then use Amazon SNS to alert your team

B. Set up a cron job to actively monitor the AWS CloudTrail logs for increased traffic and use Amazon SNS to alert your team

C. Use an Amazon EMR job to run every thirty minutes analyze the CloudWatch logs from your application Amazon EC2 instances in a batch manner to detect a sharp increase in traffic and then use the Amazon SNS SMS notification to alert your team

D. Use an Amazon EMR job to run every thirty minutes, analyze the Elastic Load Balancing access logs in a batch manner to detect a sharp increase in traffic and then use the Amazon Simple Email Service to alert your team

E. Set up a cron job to actively monitor the AWS CloudTrail logs for increased traffic and use Amazon SNS to alert your team

Answer: A

Explanation: NetworkIn Metric: The number of bytes received by each instance on all network interfaces. The metric shows the amount of input

traffic of the network on a particular instance. The number of bytes received during the period is the number of bytes reported. This number can be divided by 300 bytes per second if you are using basic monitoring. Divide it by 60, if you have detailed monitoring.

291. Amazon SQS and Auto-scaling are used by the program to handle background work. The Auto-scaling policy is based on the amount and maximum instance count of 100 messaging in the queue. The category has never increased beyond 50 since the application was launched. The Auto-scaling group is now 100, the queue size is growing and there are very few jobs that are completed. The number of messages sent to the queue is regular. What should you do to identify why the queue size is unusually high and reduce it?

A. Analyze the application logs to identify possible reasons for message processing failure and resolve the cause for failures

B. Temporarily increase the AutoScaling group's desired value to 200. When the queue size has been reduced, reduce it to 50
C. Analyze CloudTrail logs for Amazon SQS to ensure that the instances Amazon EC2 role have permission to receive messages from the queue
D. Create additional Auto Scaling groups enabling the processing of the queue to be performed in parallel

Answer: A

Explanation: The best option here is to examine the application logs and fix the failure. In the application, you may have a functionality problem that is causing messages to queue up and increase the number of fleet instances within the Auto-scaling group.

292. In several Amazon EC2 instances, you have an I/O and a network-intensive application that cannot handle a large ongoing increase in traffic. Two volumes of Amazon EBS PIOPS are used in the Amazon EC2 instances, each with the identical instance.

Choose the right approach in order to reduce the load on instances with the least interference with the application.

A. Stop each instance and change each instance to a larger Amazon EC2 instance type that has enhanced networking enabled and is Amazon EBS-optimized. Ensure that RAID striping is also set up on each instance

B. Add an instance-store volume for each running Amazon EC2 instance and implement RAID striping to improve I/O performance
C. Create an AMI from each instance, and set up Auto Scaling groups with a larger instance type that has enhanced networking enabled and is Amazon EBS-optimized
D. Create an AMI from an instance, and set up an Auto Scaling group with an instance type that has enhanced networking enabled and is Amazon EBS-optimized
E. Add an Amazon EBS volume for each running Amazon EC2 instance and implement RAID striping to improve I/O performance

Answer: D

Explanation: An Amazon Machine Image (AMI) gives the necessary information for launching an instance that is a virtual cloud-based server. When you launch an instance, you specify an AMI, and you can start as many instances as you need from an AMI. You can also launch instances from as many AMIs as you want.

293. You developed a new feature using AWS services. You want to test it from inside a staging VPC. How would you do this with the fastest turnaround time?

A. Use an Amazon EC2 instance that frequently polls the version control system to detect the new feature, use AWS CloudFormation and Amazon EC2 user data to run any testing harnesses to verify application functionality and then use Amazon SNS to notify the development team of the results

B. Use an Elastic Beanstalk application that polls the version control system to detect the new feature, use AWS CloudFormation and Amazon EC2 user data to run any testing

harnesses to verify application functionality and then use Amazon Kinesis to notify the development team of the results

C. Launch an Amazon Elastic Compute Cloud (EC2) instance in the staging VPC in response to a development request, and use configuration management to set up the application. Run any testing harnesses to verify application functionality and then use Amazon Simple Notification Service (SNS) to notify the development team of the results

D. Use AWS CloudFormation to launch an Amazon EC2 instance using Amazon EC2 user data to run any testing harnesses to verify application functionality and then use Amazon Kinesis to notify the development team of the results

Answer: C

Explanation: It would take more time to install Amazon Kinesis and would not be ideal to notify the concerned team as shortly as possible.

Since the test must be performed at the staging VPC, it is best to launch EC2 in the staging VPC.

The best answer to this question would be the management of AWS configuration together with SNS.

AWS Config provides a detailed inventory of current AWS resources and records configuration modifications on a continuous basis such as the tags value in the instance, security group entry/exit rules, and network ACL rules in VPCs (see the AWS Config website for the list of supported AWS resources). The AWS Config allows customers to determine how a resource was configured at any time, to view resource dependencies and to send notifications when the resource settings change. The AWS Config Rules are a new package that enables customers to assess whether their AWS resources meet the configuration requirements. In order to assess compliance of AWS resources, customers can either use predefined AWS-managed rules or define themselves.

The application should be tested in a staging VPC that is not described in option A, therefore, option C is correct.

294. Your company assigns you the management of application that uses Amazon SDK and Amazon EC2 roles to store and retrieve Amazon S3 data, access multiple tables of DynamoDB and exchange

messages using Amazon SQS queues. Your Compliance Vice-President is concerned that your security practices is not outstanding. He asked you to check that the application AWS access key is not older than six months, and to provide control evidence that these keys are rotated at least once every six months. Which option suits the best to provide the required information to VP?

A. Update your application to log changes to its AWS access key credential file and use a periodic Amazon EMR job to create a compliance report for your VP

B. Create a script to query the IAM list-access keys API to get your application access key creation date and create a batch process to periodically create a compliance report for your VP
C. Create a new set of instructions for your configuration management tool that will periodically create and rotate the application's existing access keys and provide a compliance report to your VP
D. Provide your VP with a link to IAM AWS documentation to address the VP's key rotation concerns

Answer: B

Explanation: To know when access keys have been developed, use the "iam: ListAccessKeys". This knowledge can be used to identify which keys are older than six months. Execute a batch process to generate a conformity report as requested by the Department of Compliance VP.

295. You have an application that has mandate requirements for security and compliance and the protected health information that belongs to your application should be encrypted both at rest and transit. The data flows through the load balancer and is stored on Amazon EBS volumes using three-architecture for processing. The outputs are stored in S3 using AWS SDK service. Choose the two options which allow fulfilling the security requirements.

A. Use SSL termination with a SAN SSL certificate on the load balancer. Amazon EC2 with all Amazon EBS volumes using Amazon EBS encryption, and Amazon S3 with server-side

encryption with customer-managed keys

B. Use TCP load balancing on the load balancer. SSL termination on the Amazon EC2 instances. OS-level disk encryption on the Amazon EBS volumes and Amazon S3 with server-side encryption
C. Use SSL termination on the load balancer, Amazon EBS encryption on Amazon EC2 instances and Amazon S3 with server-side encryption
D. Use TCP load balancing on the load balancer. SSL termination on the Amazon EC2 instances and Amazon S3 with server-side encryption
E. Use SSL termination on the load balancer an SSL listener on the Amazon EC2 instances, Amazon EBS encryption on EBS volumes containing PHI and Amazon S3 with server-side encryption

Answer: B and E

Explanation:

HTTPS/SSL Listeners

The following security features can be used to create a load balancer:

SSL Server Certificates

You have to deploy X.509 certificates (SSL server certificates) on your load balancer if you use HTTPS or SSL for your front-end connections. Before sending requests to the backend instance (known as SSL termination), the load balancer decodes requests from clients.

You can use TCP for front and back-end connections and deploy certificates on registered instances processing requests if you do not want a load balancer to handle SSL termination (known as SSL offloading).

Create a classic load balancer with an HTTPS Listener

A load balancer receives requests from customers and distributes the load balancer requests throughout the EC2 instances registered with the load balancer.

You can create a load balancer that listens on both the HTTP (80) and HTTPS (443) ports. If you specify the HTTPS listener to send requests to port 80 instances, the load balancer ends the requests, and no load balancer communications to the instances are encrypted. If the HTTPS listener sends

requests to port 443 instances, the load balancer communication is encrypted to the instances.

Options A & C is not correct because the transit between ELB and EC2 instances is missing in encryption.

Option D is incorrect because the data related to the EC2 instances lack encryption at rest.

296. Ace has an organization that has an AWS application with three tiers: frontend, backend and database. Different AWS services including EC2, ELB, Auto-scaling, Route53, RDS, and others are being used. The RTO (Recovery Time Objective) is set at 1 hour for the entire application. What can help you achieve this goal?

A. Create a Jenkins pipeline to automatically create AMIs for EC2 instances. Execute the pipeline every hour

B. Create regular EBS snapshots every hour using EBS lifecycle manager
C. Create a warm standby in another region. Use Route53 failover routing policy to route to the standby if the active application has an outage
D. In RDS, configure each database to create regular automated snapshots every hour. Copy the snapshots to another region

Answer: C

Explanation: The question asks for the strategy that can help the most to achieve the goal. The program includes multiple elements and all of them should be considered. The RTO (Recovery Time Objective) is set as 1 hour, which means the application will recover within 1 hour after the failure. The RPO (Recovery Point Objective) does not mention in the issue. It is a standard approach to trying to accomplish a quick RTO. It can be the recovery time because the rest of the service is running in Warm standby.

297. During a deployment cycle, you recently found a major error in your web application. It took four hours by the team during this unsuccessful deployment to return to a previously functional state, which left customers with poor user experience. Your team discussed the need to roll back failed deployments quicker and more robust.

Your web application is running on Amazon EC2 and is using Elastic Load Balancing to balance your load. How do you solve the problem?

A. Using Elastic Beanstalk, redeploy your web application and use the Elastic Beanstalk API to trigger a FailedDeployment API call to initiate a rollback to the previous version

B. Create deployable versioned bundles of your application. Store the bundle on Amazon S3. Use an AWS OpsWorks stack to redeploy your web application and use AWS OpsWorks application versioning to initiate a rollback during failures
C. Use an AWS OpsWorks stack to re-deploy your web application and use AWS OpsWorks DeploymentCommand to initiate a rollback during failures
D. Create deployable versioned bundles of your application. Store the bundle on Amazon S3. Re-deploy your web application on Elastic Beanstalk and enable the Elastic Beanstalk auto–rollback feature tied to CloudWatch metrics that define failure

Answer: C

Explanation: AWS DeploymentCommand contains a rollback option in it. Apps can be used by the following commands:

deploy: Deploy App.

Ruby on Rails apps has an optional migrate args parameter. To migrate the database, set Args to {"migrate":["true"]}.

The default setting is {"migrate":["false"]}.

The app will roll back to the previous version with the "rollback" feature.

AWS OpsWorks stores the previous versions, up to five versions, when we update an app.

We can roll an app back in four versions with this command.

298. Addis has a Chef Version 11.10 running on AWS OpsWorks Stack. He has its own cookbook hosted on Amazon S3, and this is specified in the stack as a custom cookbook. A cookbook located in an external Git repository is required, which is an open source cookbook. How could he use both of the custom books?

A. In the AWS OpsWorks stack settings, enable Berkshelf. Create a new cookbook with a Berksfile that specifies the other two cookbooks. Configure the stack to use this new cookbook
B. In the OpsWorks stack settings, add the open source project's cookbook details in addition to your cookbook

C. In your cookbook, create an S3 symlink object that points to the open source project's cookbook
D. Contact the open source project's maintainers and request that they pull your cookbook into theirs. Update the stack to use their cookbook

Answer: A

Explanation: You need a way to install and manage dependencies to use an external cookbook in an instance. A cookbook that supports a dependency manager named Berkshelf is the preferred approach. In addition to work with the test kitchens and Vagrants, Berkshelf is working on Amazon EC2 instances such as AWS OpsWorks Stacks instances.

299. A group of developers within your company wishes to move its current application in Elastic Beanstalk and use Elastic Load Balancing and Amazon SQS. Currently, you have a custom application server. How can this requirement be achieved?

A. Configure an AWS OpsWorks stack that installs the third-party application server and creates a load balancer and an Amazon SQS queue and then deploys it to Elastic Beanstalk

B. Create a custom Elastic Beanstalk platform that contains the third-party application server and runs a script that creates a load balancer and an Amazon SQS queue
C. Configure an Elastic Beanstalk platform using AWS OpsWorks deploy it to Elastic Beanstalk and run a script that creates a load balancer and an Amazon SQS queue
D. Use a Docker container that has the third-party application server installed on it and create an application source bundle that produces the load balancer and an Amazon SQS queue

Answer: D

Explanation: Elastic Beanstalk allows the use of Docker server for web applications. You can set up your own runtime environment with Docker containers. You can choose your own platform, programming language, and applications that are not supported by any other platforms (such as package managers or tools). Docker containers are autonomous and hold all configuration and software details that your web application needs to work with.

300. Your web application is running on three M3 instances in three AZs. A group of three to thirty instances are scaled using the Auto-scaling group. When you review the CloudWatch metrics, you see that there are sometimes 15 instances in your Auto Scaling group. The web application reads and writes to a DynamoDB-configured backend with 800 Write and read capacity units. The company's ID is your DynamoDB main key. In your web application, you receive 25 TB of data. You have one customer who complains delay in load time when his employees arrive at the office at 9:00 am and load the website consisting of content drawn out by DynamoDB. Other customers use the web application routinely. Select the response to ensure high availability and reduce access times for the customer.

A. Double the number of Read Capacity Units in your DynamoDB instance because the instance is probably being throttled when the customer accesses the website and your web application

B. Add a caching layer in front of your web application by choosing ElastiCache Memcached instances in one of the AZs
C. Implement an Amazon SQS queue between your DynamoDB database layer and the web application layer to minimize the large burst in traffic the customer generates when everyone arrives at the office at 9:00 am and begins accessing the website
D. Change your Auto Scaling group configuration to use Amazon C3 instance types, because the web application layer is probably running out of compute capacity
E. Use data pipelines to migrate your DynamoDB table to a new DynamoDB table with a primary key that is evenly distributed

across your dataset. Update your web application to request data from the new table

Answer: E

Explanation: The table's provisioned throughput optimal depends on the following factors:

- The primary key selection
- The workload patterns on individual items

Each item in a table is uniquely defined by the primary key. The primary key can be simple (partition key) or composite (partition key and sort key).

DynamoDB divides the items of a table into multiple partitions when saving data and mainly distributes the data based on the key partition value. Consequently, you keep the workload evenly across key partition values to achieve the full amount of query throughput provided for a table. Distributing requests across partition key values distributes the requests across partitions.

We can create a new index when we import data from S3 with the datapipeline into a new dynamodb table.

Following are the steps:

i. Login to the console and select DynamoDB.

ii. Select the table you have to copy.

iii. Select Export / Import. For copying DynamoDB table to S3 or from S3 to DynamoDB table export / import uses dataPipeline and EMR.

iv. If you do not have two IAM roles for export / import, you must create them.

v. Click the export table option.

vi. You will need to specify the following:
    - S3 bucket to copy the table data and another bucket to store log files for operation. The same bucket can be used
    - The percentage of the throughput capacity for the table to be used to read data from the table (to copy to

S3) that was provided. The default value is 25%. The increased percentage will accelerate backup

- The IAM roles: The values will be default

vii. Choose the option "create data pipeline", and the backup will be scheduled depending on the table size, the backup may take time.

viii. After exporting, check logs to verify that no bugs are there.

ix. Note the hash and range key information of the table.

x. Delete the table.

xi. Create a table with the right index. Set the provisioned throughput.

xii. Using import option, import into a table from S3.

xiii. After completion, do check that it is error free.

301. You want to allow automated testing of your CloudFormation template as part of your deployment pipeline. What assessments need to be carried out to ensure that feedback is quicker and costs and risk are minimized? (Choose 3)

A. In the testing environment, while creating the stack, specify an Amazon SNS topic for subscription. Your testing environment runs tests when it receives notification when the stack is created or updated

B. In the testing environment, create and update the stack with the template. If the template fails rollback will return the stack and its resources to exactly the same state

C. Validate the AWS CloudFormation template against the official XSD scheme definition published by Amazon Web Services

D. Use the AWS CloudFormation Validate Template to validate the syntax of the template

E. Validate the template’s is syntax using a general JSON parser

F. Use the AWS CloudFormation Validate Template to validate

the properties of resources defined in the template

Answer: A, B, and D

Explanation: The command AWS CloudFormation validate-template is designed to test the template syntax only. It does not guarantee that the property values for that resource that you listed are correct. Nor does it determine the number of services that will be available when the stack is created. You can attempt to create the stack to test the operational validity. AWS CloudFormation stack does not include a sandbox or trial area, so you will be paid for the resources you build during your study.

302. An application is written in Go Programming language, which has to be deployed to AWS. The application code is stored on a Git repository. Choose two methods for this deployment.

A. Write a Dockerfile that installs the Go base image and fetches your application using Git. Create an AWS CloudFormation template that creates and associates an AWS::EC2::Instance resource type with an AWS::EC2::Container resource type
B. Write a Dockerfile that installs the Go base image and uses Git to fetch your application. Create a new AWS OpsWorks stack that contains a Docker layer that uses the Dockerrun.aws.json file to deploy your container and then use the Dockerfile to automate the deployment
C. Create a new AWS Elastic Beanstalk application and configure a Go environment to host your application. Using Git, check out the latest version of the code, once the local repository for Elastic Beanstalk is configured using "eb create" command to create an environment and then use "eb deploy" command to deploy the application
D. Write a Dockerfile that installs the Go base image and fetches your application using Git. Create a new AWS Elastic Beanstalk application and use this Dockerfile to automate the deployment

Answer: C and D

Explanation: Option B is incorrect because OpsWorks works with Chef recipes and not with Docker containers.

Option A is incorrect because there is no AWS::EC2::Container resource for CloudFormation.

The deployment of a web application from Docker containers is supported by Elastic Beanstalk. You can set your own runtime environment with Docker containers. You can set your own runtime environment with Docker containers. You can choose your own platform, the language of programming, and any application dependencies that are not supported by other platforms, such as package managers or tools. Docker containers are independent and contain all configuration information and software required to execute your web application.

303. Adrian has multi-level architecture that uses a load balancer and is supported by a web tier within an Amazon EC2 Auto-scaling group. This architecture serves public facing web traffic. During a traffic peak, Adrian notices that the amount of incoming traffic and Auto-scaling policy that he had setup is adding new instances disproportionately. How should he stop the Auto-scaling group in reaction to incoming traffic from scaling incorrectly?

A. Using a custom CloudWatch metric, insert the elapsed time since the instance launch to the time the instance responds to an Elastic Load Balancing health check and periodically adjust the Pause Time of the UpdatePolicy and reduce the Scaling Adjustment property by 50%

B. Using CloudWatch and the instance BootTime metric, increase the PauseTime and CoolDown property on the Auto Scaling group to be over the value of the metric

C. Using a custom CloudWatch metric, insert the elapsed time since the instance launch to the time the instance responds to an Elastic Load Balancing health check, and periodically adjust the Pause Time and the CoolDown property on the AutoScaling group to be over the value of the metric

D. Using a third-party configuration management tool and the AWS SDK, suspend all ScheduledActions of the Auto-scaling group until after the traffic peak and then resume all scheduledActions

Answer: C

Explanation: The question focuses on adding traffic-related instances with an Auto-scaling group.

To control how rolling updates are performed when changes are made to the launch configuration of the auto scaling group, you can add a UpdatePolicy attribute to your Auto-scaling group. It is mainly used in combination with AutoScalingGroup resource CloudFormation Templates.

In AutoScalingGroup's UpdatePolicy attribute, PauseTime is used.

If you do not have the correct settings configured, a rolling update on an Auto-scaling group can lead to unexpected behaviors.

PauseTime refers to the amount of time that AWS CloudFormation takes to start Software applications after making a change to a set of instances. For example, to scale up the instances in an Auto-scaling group, you might need to specify PauseTime.

304. You have created a web application based on the Auto-scaling group of web servers using ELB. You create a new AMI for every application version deployment. You are releasing a new version of the application in which you want a controlled migration of users with the A/B deployment technique, while the fleet is constant 12 hours long to ensure that the new version is functioning. What method should you opt for to enable this technique while easily rolling back when required?

A. Create an Auto-scaling launch configuration with the new AMI. Configure Auto-scaling to vary the proportion of instances launched from the two launch configurations
B. Create a load balancer. Create an Auto-scaling launch configuration with the new AMI to use the new launch configuration and to register instances with the new load balancer. Use Amazon Route53 weighted Round Robin to vary the proportion of requests sent to the load balancers
C. Launch new instances using the new AMI and attach them to the Auto-scaling group. Configure Elastic Load Balancing to vary the proportion of requests sent to instances running the two application versions
D. Create an Auto-scaling launch configuration with the new

AMI. Configure the Auto-scaling group with the new launch configuration. Use the Auto-scaling rolling updates feature to migrate to the new version

E. Create an Auto-scaling launch configuration with the new AMI. Create an Auto-scaling group configured to use the new launch configuration and to register instances with the same load balancer. Vary the desired capacity of each group to migrate

Answer: B

Explanation: Because you want to control the use of the new application, the best way is to use the weighted method of Route53. Weighted routing can assign a single domain (example.com) or subdomain name (acme.example.com) to multiple resources and choose the amount of traffic that is routed to each resource. This can be helpful for a range of purposes such as load balancing and testing of new software versions.

305. Your Company uses a third-party configuration management tool for web application development environment in order to create a Docker container on a local developer's machine. What are you supposed to do for making sure that your application is not effected by the web application, supporting network storage and security infrastructure after the deployment for staging and production environments in AWS occurs?

A. Define an AWS CloudFormation template to place your infrastructure into version control and use the same template to deploy the Docker container into Elastic Beanstalk for staging and production
B. Because the application is inside a Docker container, there are no infrastructure differences to be taken into account when moving from the local development environments to AWS for staging and production
C. Write a script using the AWS SDK or CLI to deploy the application code from version control to the local development environments staging and production using AWS OpsWorks
D. Define an AWS CloudFormation template for each stage of the application deployment lifecycle–development, staging, and

production – and have tagged in each template to define the environment.

Answer: A

Explanation: The deployment of a web application from Docker containers is supported by Elastic Beanstalk. You can set your own runtime environment with Docker containers. You can choose a platform, programming language, and any application dependencies that are not supported by other platforms, such as package managers or tools. Docker containers are autonomous and contain all the configuration information and software needed to run your web application.

If you create your infrastructure by using Docker with Elastic Beanstalk, it handles capacity supply details, load-balancing, scaling and application health monitoring automatically.

306. Your company has set up automated resources on AWS and you are assigned to take care of these resources. Your task is to integrate some of the company's chef recipes to be used for the existing OpsWorks stack in AWS. The problem is, when you visited the recipes section, you could not find the option to add any recipes. What could be reason for this and how will you fix it?

A. The stack layers were created without the custom cookbooks option. Change the layer setting accordingly

B. The stack layers were created without the custom cookbooks option. Just change the stack settings accordingly
C. Once you create a stack, you cannot assign custom recipes, this needs to be done when the stack is created
D. Once you create layers in the stack, you cannot assign custom recipes, this needs to be done when the layers are created

Answer: B

Explanation: This is mentioned in the AWS documentation: To have a stack install and use custom cookbooks, you must configure the stack to enable custom cookbooks. If it is not already configured, you must then provide the repository URL and any related information such as password.
https://docs.aws.amazon.com/opsworks/latest/userguide/workingcookbook-

installingcustom-enable.html

307. You are a DevOps engineer for a company that has a number of CloudFormation templates in AWS. The IT security department wants to know all the users who use CloudFormation stacks in the company's AWS account. Which of the following can be done to take care of this security concern?

A. Connecting SQS and CloudFormation so that a message is published for each resource created in the CloudFormation stack

B. Enabling CloudWatch logs for each CloudFormation stack to track the resource creation events
C. Enabling CloudTrail logs so that the API calls can be recorded
D. Enabling CloudWatch events for each CloudFormation stack to track resource creation events

Answer: C

Explanation: This is given as a best practice on the AWS documentation: AWS CloudTrail tracks anyone making AWS CloudFormation API calls in your AWS account. API calls are logged whenever anyone uses the AWS CloudFormation API, the AWS CloudFormation console, a back-end console, or AWS CloudFormation AWS CLI commands. Enable logging and specify an Amazon S3 bucket to store the logs. That way, if you ever need to, you can audit who made what AWS CloudFormation call in your account.
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html

308. Your development team is planning for continuous release cycles for their application. They want to use the AWS service available to be able to deploy a web application and rollback to its previous version quickly if required. Which of the following can be used to meet this requirement? (Choose 2)

A. Use the CloudFormation service. Create separate templates for each application revision and deploy them accordingly

B. Use the OpsWorks service to deploy the web instances. Deploy the app to the OpsWorks web layer. Rollback using the Deploy app in OpsWorks
C. Use the Elastic BeanStalk service. Create separate environments for each application revision. Revert back to an environment in case the new environment does not work
D. Use the Elastic beanstalk service. Use Application versions and upload the revisions of your application. Deploy the revisions accordingly and rollback to prior versions accordingly

Answer: B and D

Explanation: The AWS documentation states:

In Elastic BeanStalk, and application version refers to a specific, labeled iteration of deployable code for a web application. An application version points to an Amazon Simple Storage Service (Amazon S3) object that contains the deployable code such as a Java WAR file. An application version is part of an application. Applications can have many versions and each application version is unique. In a running environment, you can deploy any application version you already uploaded to the application or you can upload and immediately deploy a new application version. You might upload multiple application versions to test differences between one version of your web application to another.
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/concepts.html

AWS OpsWorks Stacks app represents code that you want to run on an application server. The code itself resides in a repository such as an Amazon S3 archive; the app contains the information required to deploy the code to the appropriate application server instances.
https://docs.aws.amazon.com/opsworks/latest/userguide/workingapps.html

309. On AWS, to be able to scale out to react the increase in demand quickly, you need to build out a layer in a software stack. You are running the code on EC2 instances in an Auto-scaling group behind an ELB. Which application deployment should you use?

A. Create a new Auto Scaling Launch Configuration with UserData scripts configured to pull the latest code at all times

B. Create a Dockerfile when preparing to deploy a new version to production and publish it to S3. Use UserData in the Auto Scaling Launch configuration to pull down the Dockerfile from S3 and run it when new instances launch
C. Bake an AMI when deploying new versions of code, and use that AMI for the Auto Scaling Launch Configuration
D. SSH into new instances that come online, and deploy new code onto the system by pulling it from an S3 bucket, which is populated by code that you refresh from source control on new pushes

Answer: C

Explanation: Since the time required to spin up an instance is required to be fast, it is better to create an AMI rather than use User Data. When you use User Data, the script will be run during boot up, and hence this will be slower. An Amazon Machine Image (AMI) provides the information required to launch an instance, which is a virtual server in the cloud. You specify an AMI when you launch an instance, and you can launch as many instances from the AMI as you need. You can also launch instances from as many different AMIs as you need.

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AMIs.html

310. To coordinate the creation of stack resources in CloudFormation, which of the following can be used? (Choose 2)

A. CreationPolicy attribute

B. HoldPolicy attribute
C. AWS::CloudFormation::WaitCondition
D. AWS::CloudFormation::HoldCondition

Answer: A and C

Explanation: The AWS documentation mentions that using the AWS::CloudFormation::WaitCondition resource and CreationPolicy attribute, you can do the following:

- Coordinate stack resource creation with other configuration actions that are external to the stack creation

- Track the status of a configuration process

https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-waitcondition.html

311. How much time does the OpsWorks service waits for a response before deeming an underlying instance as a failed instance?

A. 60 minutes

B. 20 minutes
C. 5 minutes
D. 1 minute

Answer: C

Explanation: As per the AWS documentation, every instance has an AWS OpsWorks Stacks agent that communicates regularly with the service. AWS OpsWorks Stacks uses that communication to monitor instance health. If an agent does not communicate with the service for more than approximately five minutes, AWS OpsWorks Stacks considers the instance to have failed.
https://docs.aws.amazon.com/opsworks/latest/userguide/workinginstances-autohealing.html

312. You are designing an OpsWorks stack in AWS. Your company's on-premises Chef configuration has some custom recipes that are required to run whenever an instance is launched in OpsWorks. Which steps should be carried out to fulfill this requirement? (Choose 2)

A. Ensure the recipe is placed as part of the Setup Lifecycle event as part of the Stack setting

B. Ensure the recipe is placed as part of the Setup Lifecycle event as part of the Layer setting
C. Ensure the custom cookbooks option is set in OpsWorks layer
D. Ensure the custom cookbooks option is set in OpsWorks stack

Answer: B and D

Explanation: Each layer has a set of built-in recipes assigned to each

lifecycle event, although some layers lack Undeploy recipes. When a lifecycle event occurs on an instance, AWS OpsWorks Stacks runs the appropriate set of recipes for the associated layer. For more information on automating recipes, follow the below URL:
https://docs.aws.amazon.com/opsworks/latest/userguide/workingcookbook-assigningcustom.html

313. From the following options, which one is a reliable and durable logging solution to track changes made to your AWS account?

A. Create three new CloudTrail trails with three new S3 buckets to store the logs one for the AWS Management console, one for AWS SDKs and one for command line tools. Use IAM roles and S3 bucket policies on the S3 buckets that store your logs

B. Create a new CloudTrail trail with an existing S3 bucket to store the logs and with the global services option selected. Use S3 ACLs and Multi Factor Authentication (MFA) Delete on the S3 bucket that stores your logs
C. Create a new CloudTrail with one new S3 bucket to store the logs. Configure SNS to send log file delivery notifications to your management system. Use IAM roles and S3 bucket policies on the S3 bucket that stores your logs
D. Create a new CloudTrail trail with one new S3 bucket to store the logs and with the global services option selected. Use IAM roles S3 bucket policies and Multi Factor Authentication (MFA) Delete on the S3 bucket that stores your logs

Answer: D

Explanation: AWS Identity and Access Management (IAM) is integrated with AWS CloudTrail, a service that logs AWS events made by or on behalf of your AWS account. CloudTrail logs authenticated AWS API calls and also AWS sign-in events, and collects this event information in files that are delivered to Amazon S3 buckets. You need to ensure that all services are included. Hence option B is partially correct.
https://docs.aws.amazon.com/IAM/latest/UserGuide/cloudtrail-integration.html

314. To create the current application in a specified environment, which of the following commands for the Elastic BeanStalk CLI can be used?

A. en app

B. en env
C. eb start
D. eb create

Answer: D

Explanation: EB is a Command Line Interface (CLI) tool for Elastic BeanStalk that you can use to deploy your applications quickly and more easily. EB CLI introduces the commands eb create, eb deploy, eb open, eb console, eb scale, eb setenv, eb config, eb terminate, eb clone, eb list, eb use, eb printenv, and eb ssh. In EB CLI 3.1 or later, you can also use the eb swap command. In EB CLI 3.2 only, you can use the eb abort, eb platform, and eb upgrade commands. In addition to these new commands, EB CLI 3 commands differ from EB CLI 2.6 commands in several cases:

For more information on EB CLI, follow the below URL:

https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/eb-cli.html#eb-cli2-differences

315. Your company is planning to automate its integration, build, and deployment processes by using available AWS services. It is a part of plan to use AWS CodeBuild to build the artifacts. When using AWS CodeBuild, which of the following files is used to specify a collection of build commands that can be used by the service during the build process?

A. appspec.json

B. buildspec.xml
C. buildspec.yml
D. appspec.yml

Answer: C

Explanation: It is mentioned in the AWS documentation: AWS CodeBuild currently supports building from the following source code repository providers. The source code must contain a build specification (build spec) file, or the build spec must be declared as part of a build project definition. A ***buildspec*** is a collection of build commands and related settings, in YAML format, that AWS CodeBuild uses to run a build.
https://docs.aws.amazon.com/codebuild/latest/userguide/planning.html

316. As a DevOps engineer, it is your task to create a continuous integrated and continuous delivery model for your organization's application. Which of the following services are best suited for this purpose? (Choose 2)

A. AWS SQS

B. AWS IAM
C. AWS CodeDeploy
D. AWS Code PipeLine

Answer: C and D
Explanation: AWS CodeDeploy is a deployment service that automates application deployments to Amazon EC2 instances or on-premises instances in your own facility. You can deploy a nearly unlimited variety of application content, such as code, web and configuration files, executables, packages, scripts, multimedia files, and so on. AWS CodeDeploy can deploy application content stored in Amazon S3 buckets, GitHub repositories, or Bitbucket repositories.
AWS CodePipeline is a continuous delivery service you can use to model, visualize, and automate the steps required to release your software. You can quickly model and configure the different stages of a software release process. AWS CodePipeline automates the steps required to release your software changes continuously.
You can learn more about CodeDeploy and Code PipeLine using the given below URLs:
https://docs.aws.amazon.com/codedeploy/latest/userguide/welcome.html
https://docs.aws.amazon.com/codepipeline/latest/userguide/welcome.html

317. To make sure that the EC2 instances can work with AWS

CodeDeploy, which of the following pre-requisites must be ensured? (Choose 2)

A. The CodeDeoploy agent is installed on the EC2 instance

B. The EC2 instance is placed in the default VPC
C. The EC2 instance is configured with Enhanced Networking
D. An IAM role is attached to the instance so that it can work with the CodeDeploy service

Answer: A and **D**
Explanation: These pre-requisites are clearly mentioned in the AWS documentation. You can learn more about instance for CodeDeploy by visiting the given URL:
https://docs.aws.amazon.com/codedeploy/latest/userguide/instances.html

318. You have production instances running in an Auto-scaling group and it is decided that you need to change the instance type of these instances. You currently have four instances and you cannot afford interruption in service and you want the surety that two instances are running during the update. Which of the below options can be used for this?

A. AutoScalingIntegrationUpdate

B. AutoScalingReplacingUpdate
C. AutoScalingRescheduledAction
D. AutoScalingRollingUpdate

Answer: D
Explanation: The AWS::AutoScaling::AutoScalingGroup resource supports an UpdatePolicy attribute. This is used to define how an Auto-scaling group resource is updated when an update to the CloudFormation stack occurs. A common approach to updating an Auto Scaling group is to perform a rolling update, which is done by specifying the AutoScalingRollingUpdate policy. This retains the same Auto-scaling group and replaces old instances with new ones, according to the parameters specified. For more information on Auto-scaling updates, please refer to the given URL:

https://aws.amazon.com/premiumsupport/knowledge-center/auto-scaling-group-rolling-updates/

319. As a DevOps engineer, you are instructed to deploy Docker containers using AWS OpsWorks. How will you do this? (Choose 2)

A. Use CloudFormation to deploy Docker containers since this is not possible in OpsWorks. Then attach the CloudFormation resources as a layer in OpsWorks

B. Use custom cookbooks for your OpsWorks stack and provide the Git repository which has the chef recipes for the Docker containers
C. In the App for OpsWorks deployment, specify the Git URL for the recipes which will deploy the applications in the Docker environment
D. Use Elastic beanstalk to deploy Docker containers since this is not possible in OpsWorks. Then attach the elastic BeanStalk environment as a layer in OpsWorks

Answer: B and C

Explanation: AWS OpsWorks lets you deploy and manage application of all shapes and sizes. OpsWorks layers let you create blueprints for EC2 instances to install and configure any software that you want. You can refer to the given URL for more information about Docker and OpsWorks.
https://aws.amazon.com/blogs/devops/running-docker-on-aws-opsworks/

320. For automatic collection of software inventory and applying OS patches on EC2 instances, which of the following AWS tools can be helpful?

A. EC2 AMIs

B. AWS Code PipeLine
C. EC2 Systems Manager
D. AWS CodeDeploy

Answer: C

Explanation: The Amazon EC2 Systems Manager helps you automatically collect software inventory, apply OS patches, create system images, and configure Windows and Linux operating systems. These capabilities enable automated configuration and ongoing management of systems at scale, and help maintain software compliance for instances running in Amazon EC2 or on-premises. One feature within Systems Manager is Automation, which can be used to patch, update agents, or bake applications into an Amazon Machine Image (AMI). With Automation, you can avoid the time and effort associated with manual image updates, and instead build AMIs through a streamlined, repeatable, and auditable process.

https://aws.amazon.com/blogs/aws/streamline-ami-maintenance-and-patching-using-amazon-ec2-systems-manager-automation/

321. Your application uses the EC2/On-premises compute platform. Which of the following files needs to be included along with source code binaries while deploying your application using the AWS CodeDeploy service?

A. appspec.yml

B. appspec.json
C. appconfig.yml
D. appconfig.json

Answer: A

Explanation: The AWS documentation mentions:

The application specification file (AppSpec file) is a YAML-formatted file used by AWS CodeDeploy to determine:

- What it should install onto your instances from your application revision in Amazon S3 or GitHub
- Which lifecycle event hooks to run in response to deployment lifecycle events

An AppSpec file must be named appspec.yml and it must be placed in the root of an application's source code's directory structure. Otherwise, deployments will fail.

https://docs.aws.amazon.com/codedeploy/latest/userguide/reference-appspec-file.html

322. Your company has a number of CloudFormation stack defined in AWS. Some of these stacks have been targeted for deletion as part of the routine housekeeping activity. Upon trying, few of those stacks are failing to delete. Which of the following could be the reason of this failure? (Choose 2)

A. The stack has an EC2 security group, which has EC2 instance attached to it

B. The stacks were created with wrong template version. Since the standard template version is now higher, it is preventing the deletion if the stacks. You need to contact AWS support
C. The stack has an S3 bucket defined, which has objects present in it
D. The stack consists of an EC2 resource, which was created with a custom AMI

Answer: A and C
Explanation: It is mentioned in the AWS documentation:
Some resources must be empty before they can be deleted. For example, you must delete all objects in an Amazon S3 bucket or remove all instances in an Amazon EC2 security group before you can delete the bucket or security group.
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/troubles

323. Which of the following are the offerings of AWS to secure data at rest and in transit? (Choose 3)

A. Using IOPS volumes when working with EBS volumes on EC2 instances

B. Using server side encryption for S3
C. Encrypting all EBS volumes attached to EC2 instances
D. Using SSL/HTTPS when using the ELB

Answer: B, C, and D
Explanation: According to the AWS documentation, data protection refers to protecting data while in-transit (as it travels to and from Amazon S3) and

at rest (while it is stored on disks in Amazon S3 data centers). You can protect data in transit by using SSL or by using client-side encryption. You have the following options of protecting data at rest in Amazon S3.

- Use Server-Side Encryption – You request Amazon S3 to encrypt your object before saving it on disks in its data centers and decrypt it when you download the objects
- Use Client-Side Encryption – You can encrypt data client-side and upload the encrypted data to Amazon S3. In this case, you manage the encryption process, the encryption keys, and related tools

Amazon EBS encryption offers you a simple encryption solution for your EBS volumes without the need for you to build, maintain, and secure your own key management infrastructure. When you create an encrypted EBS volume and attach it to a supported instance type, the following types of data are encrypted:

- Data at rest inside the volume
- All data moving between the volume and the instance
- All snapshots created from the volume

You can create a load balancer that uses the SSL/TLS protocol for encrypted connections (also known as SSL offload). This feature enables traffic encryption between your load balancer and the clients that initiate HTTPS sessions, and for connections between your load balancer and your EC2 instances.

https://d0.awsstatic.com/whitepapers/aws-securing-data-at-rest-with-encryption.pdf

324. You need to create nested stacks in AWS CloudFormation. Which of the following resource should be used?

A. AWS::CloudFormation::StackNest

B. AWS::CloudFormation::NestedStack
C. AWS::CloudFormation::Nested
D. AWS::CloudFormation::Stack

Answer: D

Explanation: It is mentioned in the AWS documentation, a nested stack is a stack that you create within another stack by using the

AWS::CloudFormation::Stack resource. With nested stacks, you deploy and manage all resources from a single stack. You can use outputs from one stack in the nested stack group as inputs to another stack in the group. Refer to the below given URL for more information:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-stack-exports.html

325. Your development team uses .Net to code your web application. For continuous integration and deployment, the team wants it to be deployed to AWS. The code of application is hosted on a Git repository. Choose from the following combination of steps that needs to be taken to fulfil the requirement. (Choose 2)

A. Use a chef recipe to deploy the code and attach it to the Elastic beanstalk environment

B. Create a source bundle for the .Net code and upload it as an application revision
C. Use the Code Pipeline service to provision an IIS environment to host the application
D. Use the Elastic beanstalk service to provision an IIS platform web environment to host the application

Answer: B and D
Explanation: When you provision an environment using the Elastic BeanStalk service, you can choose the IIS platform to host the .Net based application. You can also upload the application as a zip file and specify it as an application revision.
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_NET.ht

326. You have an application hosted on AWS and you are using Jenkins as your continuous integration system. The builds are then deployed on newly launched EC2 instances. You want to minimize the overall cost of the continuous integration and deployment pipeline. Which of the below options will help in meeting the requirements? (Choose 2)

A. Ensure the instances are created beforehand for faster

turnaround time for the application builds to be placed

B. Ensure that all build tests are conducted using Jenkins before deploying the build to newly launched EC2 instances
C. Ensure that all build tests are conducted on newly launched EC2 instances
D. Ensure the instances are launched only when the build tests are completed

Answer: B and D

Explanation: To ensure low cost, one can carry out the build tests on the Jenkins server itself. Once the build tests are completed, the build can then be transferred onto newly launched EC2 Instances. Refer to the given URL for more information about AWS and Jenkins.

https://aws.amazon.com/getting-started/projects/setup-jenkins-build-server/

327. Your IT supervisor is worried about the cost incurred company's AWS resources that are hosted on AWS. The supervisor wants to monitor the cost usage. Which of the following options can be helpful in monitoring the cost of AWS resources while also looking at the possibility of cost optimization? (Choose 3)

A. Consider using the Trusted Advisor

B. Send all logs to CloudWatch logs and inspect the logs for billing details
C. Create Budgets in billing section so that budgets are set before hand
D. Use the Cost Explorer to see the costs of AWS resources

Answer: A, C, and D

Explanation: Visit the AWS Trusted Advisor console regularly. Trusted Advisor works like a customized cloud expert, analyzing your AWS environment and providing best practice recommendations to help you save money, improve system performance and reliability, and close security gaps.

Consider using budgets if you have a defined spending plan for a project or service and you want to track how close your usage and costs are to exceeding your budgeted amount. Budgets use data from Cost Explorer to

provide you with a quick way to see your usage-to-date and current estimated charges from AWS. You can also set up notifications that warn you if you exceed or are about to exceed your budgeted amount.

For a quick, high-level analysis use Cost Explorer, which is a free tool that you can use to view graphs of your AWS spend data. It includes a variety of filters and preconfigured views, as well as forecasting capabilities. Cost Explorer displays data from the last 13 months, the current month, and the forecasted costs for the next three months, and it updates this data daily.

https://aws.amazon.com/solutions/cost-optimization-monitor/

328. Your application is deployed on an EC2 instance and it needs to write data to DynamoDB table. Your security policy dictates that it is not allowed to store any security keys on the EC2 instance. Keeping the security policy and the requirement in mind, which of the following steps would you take? (Choose 2)

A. Add an IAM user to a running EC2 instance

B. Create an IAM user that allows write access to the DynamoDB table
C. Create an IAM role that allows write access to the DynamoDB table
D. Add an IAM role to a running EC2 instance

Answer: C and D

Explanation: The AWS documentation for IAM roles states, we designed IAM roles so that your applications can securely make API requests from your instances, without requiring you to manage the security credentials that the applications use. Instead of creating and distributing your AWS credentials, you can delegate permissions to make API requests using IAM roles.

http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html

329. There are multiple applications and services running in your company's production AWS account at any given time. It is required to understand where the cost is coming from. Without spending too much development time, how can you provide a good understanding

of the most cost incurring application per month?

A. Use the AWS Price API and constantly running resource inventory scripts to calculate total price based on multiplication of consumed resources over time

B. Use AWS Cost Allocation Tagging for all resources that support it. Use the Cost Explorer to analyze costs throughout the month
C. Use custom CloudWatch Metrics in your system, and put a metric data point whenever cost is incurred
D. Create an automation script that periodically creates AWS Support tickets requesting detailed intra-month information about your bill

Answer: B

Explanation: A tag is a label that you or AWS assigns to an AWS resource. Each tag consists of a *key* and a *value.* A key can have more than one value. You can use tags to organize your resources, and cost allocation tags to track your AWS costs on a detailed level. After you activate cost allocation tags, AWS uses the cost allocation tags to organize your resource costs on your cost allocation report, to make it easier for you to categorize and track your AWS costs. AWS provides two types of cost allocation tags, an *AWS generated tag* and *user defined tags.* AWS defines, creates, and applies the AWS-generated tag for you, and you define, create, and apply user-defined tags. You must activate both types of tags separately before they can appear in Cost Explorer or on a cost allocation report.
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/cost-alloc-tags.html

330. You are building an internal, for non-production use application based on Go programming language. This application used MySQL as database. Your developers are not very much familiar with the AWS environment, but you want them to be able to deploy the code with a single command line push and you want to set this up as simply as possible. Which tool is ideal for this?

A. AWS Elastic BeanStalk

B. AWS EC2 + ELB
C. AWS OpsWorks
D. AWS CloudFormation

Answer: A

Explanation: With Elastic BeanStalk, you can quickly deploy and manage applications in the AWS Cloud without worrying about the infrastructure that runs those applications. AWS Elastic BeanStalk reduces management complexity without restricting choice or control. You simply upload your application, and Elastic BeanStalk automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring. Elastic Beanstalk supports applications developed in Java, PHP, .NET, Node.js, Python, and Ruby, as well as different container types for each language.

https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/Welcome.html

331. To process messages from an SQS queue, you have a set of EC2 instances in an Auto-scaling group. The messages contain the location in S3 from where videos need to be processed by the EC2 instances. When a scale happens, it is observed that while processing a video, the EC2 instance gets terminated. How can you implement a solution to avoid this situation?

A. By increasing the maximum and minimum size for the Auto-scaling group, and changing the scaling policies so they scale less dynamically

B. By suspending the AZRebalance termination policy
C. By changing the CoolDown property for the Auto-scaling group
D. By using lifecycle hooks to ensure the processing is complete before the termination occurs

Answer: D

Explanation: This is a case where lifecycle policies can be used. The lifecycle policy can be used to put the instance in a state of Terminating:Wait, complete the processing and then send a signal to

complete the termination. Auto-scaling lifecycle hooks enable you to perform custom actions by pausing instances as Auto-scaling launches or terminates them. For example, while your newly launched instance is paused, you could install or configure software on it.
https://docs.aws.amazon.com/autoscaling/ec2/userguide/lifecycle-hooks.html

332. You have built an application that consists of a web and application server. You are going to deploy it using Elastic BeanStalk. Before the application version is deployed on the web server, it is required to run some python scripts. Which of the following can be used for this?

A. Container commands

B. Custom resources
C. Multiple Elastic BeanStalk environments
D. Docker containers

Answer: A
Explanation: The AWS documentation mentions the following; You can use the container_commands key to execute commands that affect your application source code. Container commands run after the application and web server have been set up and the application version archive has been extracted, but before the application version is deployed. Non-container commands and other customization operations are performed prior to the application source code being extracted.
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/customize-containers-ec2.html

333. Your company's application is hosted on a set of EC2 instances that are sitting behind an ELB. To host the newer version of the application, it is required to create an OpsWorks stack. You have planned to get the stack in place, carry out a level of testing, and then deploy the app at a later stage. The OpsWorks stack and layers have been setup. For testing, current ELB is being utilized but now you are noticing that your current application has stopped responding to requests. Why do you think it happened?

A. This is because the OpsWorks web layer is utilizing the current instances after the ELB was attached as an additional layer

B. The ELB would have deregistered the older instances
C. You have configured the OpsWorks stack to deploy new instances in the same domain the older instances
D. This is because the OpsWorks stack is utilizing the current instances after the ELB was attached as a layer

Answer: B

Explanation: If you choose to use an existing Elastic Load Balancing load balancer, you should first confirm that it is not being used for other purposes and has no attached instances. After the load balancer is attached to the layer, OpsWorks removes any existing instances and configures the load balancer to handle only the layer's instances. Although it is technically possible to use the Elastic Load Balancing console or API to modify a load balancer's configuration after attaching it to a layer, you should not do so; the changes will not be permanent.
https://docs.aws.amazon.com/opsworks/latest/userguide/layers-elb.html

334. You have a set of EC2 instances that was launched by an auto scaling group. These instances are sitting behind an ELB. The assurance of storing web server logs in a durable storage layer is required. This surety is required so that the logs can later be analyzed by the staff. Which of the following steps should be takes to fulfil this requirement? (Choose 2)

A. Use AWS Data Pipeline to move log data from the Amazon S3 bucket to Amazon SQS in order to process and run reports

B. On the web servers, create a scheduled task that executes a script to rotate and transmit the logs to Amazon Glacier
C. Use AWS Data Pipeline to move log data from the Amazon S3 bucket to Amazon RedShift in order to process and run reports
D. On the web servers, create a scheduled task that executes a script to rotate and transmit the logs to an Amazon S3 bucket

Answer: C and D

Explanation: Amazon S3 is the perfect option for durable storage. The AWS Documentation mentions the following on S3 Storage. Amazon Simple Storage Service (Amazon S3) makes it simple and practical to collect, store, and analyze data - regardless of format – all at massive scale. S3 is object storage built to store and retrieve any amount of data from anywhere – web sites and mobile apps, corporate applications, and data from IoT sensors or devices.

https://aws.amazon.com/s3/

Amazon Redshift is a fast, fully managed data warehouse that makes it simple and cost-effective to analyze all your data using standard SQL and your existing Business Intelligence (BI) tools. It allows you to run complex analytic queries against petabytes of structured data, using sophisticated query optimization, columnar storage on high-performance local disks, and massively parallel query execution. Most results come back in seconds.

https://aws.amazon.com/redshift/

335. You are leading a team that is responsible for an Elastic BeanStalk application. The business requires you to move to a continuous deployment model, releasing updates to the application multiple times per day with zero downtime. What will you do to ensure this with the ability to rollback immediately in a case of an emergency?

A. Create a second Elastic Beanstalk environment with the new application version, and configure the old environment to redirect clients, using the HTTP 301 response code, to the new environment

B. Develop the application to poll for a new application version in your code repository; download and install to each running Elastic Beanstalk instance

C. Create a second Elastic Beanstalk environment running the new application version, and swap the environment CNAMEs

D. Enable rolling updates in the Elastic Beanstalk environment, setting an appropriate pause time for application startup

Answer: C

Explanation: Since the requirement calls for zero downtime and for the ability roll back quickly, we need to implement Blue green deployments using the Elastic BeanStalk service. For this, we can use the SWAP URL feature is available with Elastic BeanStalk.

The AWS whitepaper on Blue green deployments mentions the following:

You can avoid this downtime by deploying the new version to a separate environment. The existing environment's configuration is copied and used to launch the green environment with the new version of the application. The new—green—environment will have its own URL. When it is time to promote the green environment to serve production traffic, you can use Elastic Beanstalk's Swap Environment URLs feature.

https://d0.awsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf

336. Your company's application consisting of web servers and AWS RDS is hosted on AWS. It is a read-heavy application and you are noticing a decrease in response time due to the load on RDS instance. To scale the database tier, which of the following measures would you take? (Choose 2)

A. Use ElastiCache in front of your Amazon RDS DB to cache common queries

B. Use SQS to cache the database queries
C. Use Auto-scaling to scale out and scale in the database tier
D. Create Amazon DB Read Replica's. Configure the application layer to query the read replicas for query needs

Answer: A and D

Explanation: According to the AWS documentation; Amazon ElastiCache is a web service that makes it easy to deploy, operate, and scale an in-memory data store or cache in the cloud. The service improves the performance of web applications by allowing you to retrieve information from fast, managed, in-memory data stores, instead of relying entirely on slower disk-based databases.

Amazon RDS Read Replicas provide enhanced performance and durability for database (DB) instances. This replication feature makes it easy to elastically scale out beyond the capacity constraints of a single DB Instance

for read-heavy database workloads. You can create one or more replicas of a given source DB Instance and serve high-volume application read traffic from multiple copies of your data, thereby increasing aggregate read throughput. Read replicas can also be promoted when needed to become standalone DB instances.

https://aws.amazon.com/elasticache/

https://aws.amazon.com/rds/details/read-replicas/

337. You are assigned with a task to automate the creation of EBS snapshots, which of the following is the best way to do this?

A. Using Cloudwatch Events to trigger the snapshots of EBS Volumes

B. Using the AWS CodeDeploy service to create a snapshot of the AWS Volumes

C. Using the AWSConfig service to create a snapshot of the AWS Volumes

D. Creating a powershell script that uses the AWS CLI to get the volumes and then running the script as a cron job

Answer: A

Explanation: The best is to use the inbuilt service from CloudWatch, as CloudWatch Events to automate the creation of EBS Snapshots. With Option D, you would be restricted to running the PowerShell script on Windows machines and maintaining the script itself. And then, you have the overhead of having a separate instance just to run that script.

The AWS Documentation mentions; Amazon CloudWatch Events delivers a near real-time stream of system events that describe changes in Amazon Web Services (AWS) resources. Using simple rules that you can quickly set up, you can match events and route them to one or more target functions or streams. CloudWatch Events becomes aware of operational changes as they occur. CloudWatch Events responds to these operational changes and takes corrective action as necessary, by sending messages to respond to the environment, activating functions, making changes, and capturing state information.

https://docs.aws.amazon.com/AmazonCloudWatch/latest/events/WhatIsCloud

338. One of your vendors who is also a user of AWS, requires access to your AWS account in a way that the vendor should be able to read protected messages in a private S3 bucket as its leisure. Which of the following is the best way to accomplish this?

A. Generate a signed S3 PUT URL and a signed S3 PUT URL, both with wildcard values and 2 year durations. Pass the URLs to the vendor

B. Create a cross-account IAM Role with permission to access the bucket, and grant permission to use the Role to the vendor AWS account
C. Create an EC2 Instance Profile on your account. Grant the associated IAM role full access to the bucket. Start an EC2 instance with this Profile and give SSH access to the instance to the vendor
D. Create an IAM User with API Access Keys. Grant the User permissions to access the bucket. Give the vendor the AWS Access Key ID and AWS Secret Access Key for the User

Answer: B

Explanation: You can use AWS Identity and Access Management (IAM) roles and AWS Security Token Service (STS) to set up cross-account access between AWS accounts. When you assume an IAM role in another AWS account to obtain cross-account access to services and resources in that account, AWS CloudTrail logs the cross-account activity.
https://aws.amazon.com/blogs/security/tag/cross-account-access/

339. Multiple development teams work in your organization on a variety of programming languages. The applications being developed by these teams have a lot of dependencies. The company has planned to move these development environments onto AWS. Which of the following is the best solution to make this move?

A. Launch separate EC2 instances to host each application type for the developer community

B. Use the OpsWorks service, crate a stack and create separate

layers for each application environment for the developer community

C. Use the Elastic BeanStalk service and use Docker containers to host each application environment for the developer community
D. Use the CloudFormation service to create Docker containers for each type of application

Answer: C

Explanation: The AWS documentation mentions the following; Elastic Beanstalk supports the deployment of web applications from Docker containers. With Docker containers, you can define your own runtime environment. You can choose your own platform, programming language, and any application dependencies (such as package managers or tools), that aren't supported by other platforms. Docker containers are self-contained and include all the configuration information and software your web application requires to run.

https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker.

340. Your company's application is hosted on a set of EC2 servers behind an ELB. These EC2 servers are launched by an Auto-scaling group. From the following options, select the example of Blue/Green deployments in AWS.

A. Use the OpsWorks service to deploy your resources. Use 2 OpsWorks layers to deploy 2 versions of your application. When the time comes for the switch, change to the alternate layer in the OpsWorks stack

B. Re-deploy your application behind a load balancer that uses Auto-scaling groups, create a new identical Auto-scaling group, and associate it to the load balancer. During deployment, set the desired number of instances on the old Auto-scaling group to zero, and when all instances have terminated, delete the old Auto-scaling group
C. Use the Elastic beanstalk service to deploy your resources. Use 2 Elastic beanstalk environments. Use Rolling deployments to

switch between the environments

D. Use a CloudFormation stack to deploy your resources. Use 2 CloudFormation stacks. Whenever you want to switch over, deploy and use the resources in the second CloudFormation stack

Answer: B

Explanation: This deployment technique is discussed in an AWS whitepaper; As you scale up the green Auto-scaling group, you can take blue Auto-scaling group instances out of service by either terminating them or putting them in Standby state. Standby is a good option because if you need to roll back to the blue environment, you only have to put your blue server instances back in service and they are ready to go. As soon as the green group is scaled up without issues, you can decommission the blue group by adjusting the group size to zero. If you need to roll back, detach the load balancer from the green group or reduce the group size of the green group to zero.

https://d0.awsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf

341. What are the wait states that occur during the scale in and scale out process when you have added lifecycle hooks to an auto scaling group? (Choose 2)

A. Terminating:Wait

B. Pending:Wait

C. Launching:Wait

D. Exitting:Wait

Answer: A and B

Explanation: As per AWS documentation; After you add lifecycle hooks to your Auto-scaling group, they work as follows:

- Auto-scaling responds to scale out events by launching instances and scale in events by terminating instances
- Auto-scaling puts the instance into a wait state (Pending:Wait orTerminating:Wait). The instance is paused until either you tell Auto-scaling to continue or the timeout period ends

https://docs.aws.amazon.com/autoscaling/ec2/userguide/lifecycle-hooks.html

342. As a DevOps engineer, you are using OpsWorks stack to rollout a collection of web servers for your company. When the instances are launched, a configuration file is needed to be set up prior to the launching of the web application hosted on these instances. To fulfil this requirement, which of the following steps would you take? (Choose 2)

A. Configure a recipe that sets the configuration file and add it to the Deploy LifeCycle Event of the specific web layer

B. Configure a recipe that sets the configuration file and add it to the Configure LifeCycle Event of the specific web layer
C. Ensure that the OpsWorks stack is changed to use custom cookbooks
D. Ensure that the OpsWorks stack is changed to use the AWS specific cookbooks

Answer: B and C

Explanation: This is mentioned in the AWS documentation:

Configure

This event occurs on all of the stack's instances when one of the following occurs:

- An instance enters or leaves the online state
- You associate an Elastic IP address with an instance or disassociate one from an instance
- You attach an Elastic Load Balancing load balancer to a layer, or detach one from a layer

For example, suppose that your stack has instances A, B, and C, and you start a new instance, D. After D has finished running its setup recipes, AWS OpsWorks Stacks triggers the Configure event on A, B, C, and D. If you subsequently stop A, AWS OpsWorks Stacks triggers the Configure event on B, C, and D. AWS OpsWorks Stacks responds to the Configure event by running each layer's Configure recipes, which update the instances' configuration to reflect the current set of online instances. The Configure event is therefore a good time to regenerate configuration files. For example,

the HAProxy Configure recipes reconfigure the load balancer to accommodate any changes in the set of online application server instances.
You can also manually trigger the Configure event by using the Configure stack command.
https://docs.aws.amazon.com/opsworks/latest/userguide/workingcookbook-events.html

343. For testing an application, the development team has requested you to create a CloudFormation template. They have requested that when the stack is deleted, the database should be preserved. How can you do this with CloudFormation?

A. In the AWS CloudFormation template, set the AWS::RDS::DBInstance's DBlnstanceClass property to be read-only

B. In the AWS CloudFormation template, set the WaitPolicy of the AWS::RDS::DBInstance's WaitPolicy property to "Retain"
C. In the AWS CloudFormation template, set the DeletionPolicy of theAWS::RDS::DBInstance's DeletionPolicy property to "Retain"
D. Ensure that the RDS is created with Read Replica's so that the Read Replica remains after the stack is torn down

Answer: C
Explanation: With the DeletionPolicy attribute you can preserve or (in some cases) backup a resource when its stack is deleted. You specify a DeletionPolicy attribute for each resource that you want to control. If a resource has no DeletionPolicy attribute, AWS CloudFormation deletes the resource by default. Note that this capability also applies to update operations that lead to resources being removed.
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-deletionpolicy.html

344. Your company has migrated to AWS. It had an Active Directory setup in the on-premises setup and wants to use the same AD for authentication. Which AWS service can help you in continuing the use of the credentials of on-premises AD for authenticating AWS

resources such as AWS WorkSpaces?

A. The ClassicLink feature on AWS

B. The Active Directory connector service on AWS
C. The AWS Simple AD service
D. The Active Directory service on AWS

Answer: B

Explanation: As per AWS documentation, AD Connector is a directory gateway with which you can redirect directory requests to your on-premises Microsoft Active Directory without caching any information in the cloud. AD Connector comes in two sizes, small and large. A small AD Connector is designed for smaller organizations of up to 500 users. A large AD Connector can support larger organizations of up to 5,000 users. https://docs.aws.amazon.com/directoryservice/latest/admin-guide/directory_ad_connector.html

345. Your log analysis mechanism takes ten days to generate a report of top ten visitors of your web application. It is now required to implement a system that can report this information in real time, ensuring that the report is always up to date, and can handle the increase in number of requests to your application. How can this requirement be fulfilled in the most cost effective way?

A. Create a multi-AZ Amazon RDS MySQL cluster, post the logging data to MySQL, and run a map reduce job to retrieve the required information on user counts

B. Configure an Auto-scaling group to increase the size of your Amazon EMR cluster
C. Post your log data to an Amazon Kinesis data stream, and subscribe your log-processing application so that is configured to process your logging data
D. Publish your log data to an Amazon S3 bucket. Use AWS CloudFormation to create an Auto-scaling group to scale your post-processing application that is configured to pull down your log files stored in Amazon S3

E. Publish your data to CloudWatch Logs, and configure your application to auto-scale to handle the load on demand

Answer: C

Explanation: The AWS documentation mentions the following; Amazon Kinesis makes it easy to collect, process, and analyze real-time, streaming data so you can get timely insights and react quickly to new information. Amazon Kinesis offers key capabilities to cost effectively process streaming data at any scale, along with the flexibility to choose the tools that best suit the requirements of your application. With Amazon Kinesis, you can ingest real-time data such as application logs, website clickstreams, IoT telemetry data, and more into your databases, data lakes and data warehouses, or build your own real-time applications using this data. Amazon Kinesis enables you to process and analyze data as it arrives and respond in real-time instead of having to wait until all your data is collected before the processing can begin.
https://aws.amazon.com/kinesis/

346. You have multiple applications hosted on AWS and it is required to store the logs from these applications in a durable storage. After three months, these logs can be moved to archval storage. Which of the following steps would you carry out to meet this requirement? (Choose 2)

A. Use Lifecycle policies to move the data onto Amazon Simple Storage service after a period of 3 months

B. Use Lifecycle policies to move the data onto Amazon Glacier after a period of 3 months
C. Store the log files as they emitted from the application on to Amazon Simple Storage service
D. Store the log files as they emitted from the application on to Amazon Glacier

Answer: B and C

Explanation: Lifecycle configuration enables you to specify the lifecycle management of objects in a bucket. The configuration is a set of one or more rules, where each rule defines an action for Amazon S3 to apply to a group of objects. These actions can be classified as follows:

- ***Transition Actions*** – In which you define when objects transition to another storage class. For example, you may choose to transition objects to the STANDARD_IA (IA, for infrequent access) storage class 30 days after creation, or archive objects to the GLACIER storage class one year after creation
- ***Expiration Actions*** – In which you specify when the objects expire. Then Amazon S3 deletes the expired objects on your behalf

https://docs.aws.amazon.com/AmazonS3/latest/dev/object-lifecycle-mgmt.html

The AWS documentation also mentions the following; Amazon Simple Storage Service (Amazon S3) makes it simple and practical to collect, store, and analyze data - regardless of format – all at massive scale. S3 is object storage built to store and retrieve any amount of data from anywhere – web sites and mobile apps, corporate applications, and data from IoT sensors or devices.

https://aws.amazon.com/s3/

347. One of your applications is running in us-west-2 region and it requires 6 EC2 instances running all the time. Which of the following deployment type will provide you 100% fault tolerance if any AZ out of the three AZs in us-west-2 becomes unavailable? (Choose 2)

A. us-west-2a with 3 instances, us-west-2b with 3 instances, us-west-2c with 3 instances

B. us-west-2a with 6 instances, us-west-2b with 6 instances, us-west-2c with 0 instances
C. us-west-2a with 4 instances, us-west-2b with 2 instances, us-west-2c with 2 instances
D. us-west-2a with 3 instances, us-west-2b with 3 instances, us-west-2c with 0 instances
E. us-west-2a with 2 instances, us-west-2b with 2 instances, us-west-2c with 2 instances

Answer: A and B

Explanation: Since you need 6 instances running all the time, only A and B can fulfill this requirement.

AWS documentation says the following about Availability Zones; When you launch an instance, you can select an Availability Zone or let us choose one for you. If you distribute your instances across multiple Availability Zones and one instance fails, you can design your application so that an instance in another Availability Zone can handle requests.
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html

348. One of your instances is reporting an unhealthy system status check. It is something you are not responsible to monitor and repair on your own. In your AWS environment, how might you automate the repair of the system status check failure?

A. Implement a third-party monitoring tool

B. Write a script that periodically shuts down and starts instances based on certain stats
C. Create CloudWatch alarms for StatuscheckFailed_System metrics and select EC2 action. Recover the instance
D. Write a script that queries the EC2 API for each instance status check

Answer: C
Explanation: Using Amazon CloudWatch alarm actions, you can create alarms that automatically stop, terminate, reboot, or recover your EC2 instances. You can use the stop or terminate actions to help you save money when you no longer need an instance to be running. You can use the reboot and recover actions to automatically reboot those instances or recover them onto new hardware if a system impairment occurs.
https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/UsingAl

349. Your company's production web servers are running on three reserved EC2 instances with EBS as root volumes. These instances have a consistent CPU load of 80%. An ELB distributes traffic to these instances. They also have production and development Multi-AZ RDS MySQL databases. This is a mission critical system. What would you recommend for cost reduction without affecting the availability?

A. Consider removing the ELB

B. Consider using spot instances
C. Consider not using Multi-AZ RDS deployment for the development database
D. Consider using on-demand instances

Answer: C

Explanation: Multi-AZ databases is better for production environments rather than for development environments, so you can reduce costs by not using this for development environments.

Amazon RDS Multi-AZ deployments provide enhanced availability and durability for Database (DB) Instances, making them a natural fit for production database workloads. When you provision a Multi-AZ DB Instance, Amazon RDS automatically creates a primary DB Instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). Each AZ runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. In case of an infrastructure failure, Amazon RDS performs an automatic failover to the standby (or to a read replica in the case of Amazon Aurora) so that you can resume database operations as soon as the failover is complete. Since the endpoint for your DB Instance remains the same after a failover, your application can resume database operation without the need for manual administrative intervention.

https://aws.amazon.com/rds/details/multi-az/

350. You have production and development instances running in your VPC. To ensure better security, you want to restrict access to production instances for the people working on development instances. Which of the following would be the best way to accomplish this using policies?

A. Define the tags on the test and production servers and add a condition to the IAM policy that allows access to specific tags

B. Launch the test and production instances in different Availability Zones and use Multi-Factor Authentication
C. Create an IAM policy with a condition, which allows access to

only instances that are used for production or development

D. Launch the test and production instances in separate VPCs and use VPC peering

Answer: A

Explanation: You can easily add tags that define which instances are production and which are development instances and then ensure these tags are used when controlling access via an IAM policy.

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html

351. Your development team is working on an application that access resources in AWS. The users of this application will be logging in via Google and Facebook. Which of the following AWS mechanisms would you imply to authenticate users with Google or Facebook?

A. Use AWS policies

B. Use Web Identity Federation

C. This is not possible

D. Modify bucket policy on website bucket to be able to access CSS bucket

Answer: B

Explanation: You can directly configure individual identity providers to access AWS resources using web identity federation. AWS currently supports authenticating users using web identity federation through several identity providers:

- Login with Amazon
- Facebook login
- Google Sign-in

https://docs.aws.amazon.com/sdk-for-javascript/v2/developer-guide/loading-browser-credentials-federated-id.html

352. As a DevOps engineer, you are given a task to log each time an instance is scaled in or scaled out from an auto scaling group. Which of the following steps would you carry out to fulfil the requirement knowing that each option forms part of the solution. (Choose 2)

A. Create a CloudWatch event that will trigger an SQS queue

B. Create an SQS queue that will write the event to CloudWatch logs
C. Create a CloudWatch event that will trigger a Lambda function
D. Write a Lambda function that will write the event to CloudWatch logs

Answer: C and D

Explanation: You can run an AWS Lambda function that logs an event whenever an Auto-scaling group launches or terminates an Amazon EC2 instance and whether the launch or terminate event was successful.
https://docs.aws.amazon.com/AmazonCloudWatch/latest/events/LogASGroup

353. An Elastic BeanStalk environment was being used by your development team. After a week, the environment was discarded and a new one was created. Upon trying to access the data on the older environment, it was not available. Why do you think it happened?

A. This is because before the environment termination, Elastic beanstalk copies the data to DynamoDB, and hence the data is not present in the EBS volumes

B. This is because the underlying EC2 Instances are created with no persistent local storage
C. This is because the underlying EC2 Instances are created with IOPS volumes and cannot be accessed once the environment has been terminated
D. This is because the underlying EC2 Instances are created with encrypted storage and cannot be accessed once the environment has been terminated

Answer: B

Explanation: According to the AWS documentation; Elastic Beanstalk applications run on Amazon EC2 instances that have no persistent local storage. When the Amazon EC2 instances terminate, the local file system is not saved, and new Amazon EC2 instances start with a default file system.

You should design your application to store data in a persistent data source.

https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/concepts.concepts.des

354. Your company has a set of EC2 instances hosted in AWS that are launched by an Auto-scaling group. The load on the servers is causing loss of requests; even though, auto scaling is launching more instances to manage the load but some requests are still getting lost. Which of the following can give you the most cost effective solution to avoid loss of recently submitted requests?

A. Pre-warm your ELB

B. Use larger instances for your application
C. Use an SQS queue to decouple the application components
D. Keep one extra EC2 instance always powered on in case a spike occurs

Answer: C

Explanation: Amazon Simple Queue Service (SQS) is a fully-managed message queuing service for reliably communicating among distributed software components and microservices - at any scale. Building applications from individual components that each perform a discrete function improves scalability and reliability, and is best practice design for modern applications.
https://aws.amazon.com/sqs/

355. Your company is concerned with EBS volume backup on EC2 and wants a proper backup solution with guaranteed durability of backup data. Which of the following solutions would you implement and why?

A. Write a cronjob that uses the AWS CLI to take a snapshot of production EBS volumes. The data is durable because EBS snapshots are stored on the Amazon S3 standard storage class

B. Use a lifecycle policy to back up EBS volumes stored on Amazon S3 for durability
C. Write a cronjob on the server that compresses the data that

needs to be backed up using gzip compression, then use AWS CLI to copy the data into an S3 bucket for durability

D. Configure Amazon Storage Gateway with EBS volumes as the data source and store the backups on premise through the storage gateway

Answer: A

Explanation: You can take snapshots of EBS volumes and to automate the process, you can use the CLI. The snapshots are automatically stored on S3 for durability.

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSSnapshots.html

356. As a DevOps engineer, one of your responsibilities is to create CloudFormation templates for your company. It is required to have an S3 bucket to store logs of all the development resources. How will you achieve this?

A. By using the metadata section in the CloudFormation template to decide on whether to create the S3 bucket or not

B. By creating an S3 bucket from before and then just providing access based on the tag value mentioned in the CloudFormation template

C. By creating a parameter in the CloudFormation template and then using the Condition clause in the template to create an S3 bucket if the parameter has a value of development

D. By creating separate CloudFormation templates for Development and production

Answer: C

Explanation: You might use conditions when you want to reuse a template that can create resources in different contexts, such as a test environment versus a production environment. In your template, you can add an EnvironmentType input parameter, which accepts either prod or test as inputs. For the production environment, you might include Amazon EC2 instances with certain capabilities; however, for the test environment, you want to use reduced capabilities to save money. With conditions, you can define which resources are created and how they are configured for each

environment type.
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/conditic section-structure.html

357. You are writing a CloudFormation template to create 2 EC2 instances behind an ELB. You want the DNS of the load balancer to be returned upon creation of the stack. Which section of the template would do this for you?

A. Mappings

B. Outputs
C. Parameters
D. Resources

Answer: B

Explanation: The below example shows a simple CloudFormation template. It creates an EC2 instance based on the AMI - ami-d6f32ab5. When the instance is created, it will output the AZ in which it is created.

```
{
	“Resources”: {
		“MyEC2Instance”: {
			"Type": "AWS::EC2::Instance",
		    "Properties": {
				"ImageId": "ami-d6f32ab5"
		    }
	}
}
	“Outputs”: {
		“Availability”: {
		“Description”: “The Instance ID”.
		“Value”:
		  { "Fn::GetAtt": [ "MyEC2Instance", "AvailabilityZone"]}
		}
		}
}
```

For more information on CloudFormation, please visit the given URL:
https://aws.amazon.com/cloudformation/

358. Your AWS infrastructure consists of EC2 instances behind an ELB and the instances are launched and terminated by an Auto-scaling group. You also have a AWS RDS MySQL database. Which of the following can be used to take you one step further towards a self-healing architecture?

A. Create one more Auto-scaling group in another region for fault tolerance

B. Create one more ELB in another region for fault tolerance
C. Enable Multi-AZ feature for the AWS RDS database
D. Enable Read Replica's for the AWS RDS database

Answer: C

Explanation: As per AWS documentation, Amazon RDS Multi-AZ deployments provide enhanced availability and durability for Database (DB) Instances, making them a natural fit for production database workloads. When you provision a Multi-AZ DB Instance, Amazon RDS automatically creates a primary DB Instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). Each AZ runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. In case of an infrastructure failure, Amazon RDS performs an automatic failover to the standby (or to a read replica in the case of Amazon Aurora), so that you can resume database operations as soon as the failover is complete. Since the endpoint for your DB Instance remains the same after a failover, your application can resume database operation without the need for manual administrative intervention.
https://aws.amazon.com/rds/details/multi-az/

359. You are storing sensitive data on AWS, which of the following steps should you take? (Choose 3)

A. Enable S3 Encryption

B. Enable EBS Encryption

C. With AWS, you do not need to worry about encryption
D. Encrypt the file system on an EBS volume using Linux tools

Answer: A, B, and D

Explanation: Data protection refers to protecting data while in-transit (as it travels to and from Amazon S3) and at rest (while it is stored on disks in Amazon S3 data centers). You can protect data in transit by using SSL or by using client-side encryption.

https://docs.aws.amazon.com/AmazonS3/latest/dev/UsingEncryption.html

Amazon EBS encryption offers you a simple encryption solution for your EBS volumes without the need for you to build, maintain, and secure your own key management infrastructure. When you create an encrypted EBS volume and attach it to a supported instance type, the following types of data are encrypted:

- Data at rest inside the volume
- All data moving between the volume and the instance
- All snapshots created from the volume

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.htm

360. The company that you work for is a startup and current funding has left it short on cash. It is not possible for your company to buy thousands of dollar of storage hardware for its application that receives huge amount of data. It has opted to use AWS. Which services would you implement to store unlimited amount of data without any effort to scale when demand unexpectedly increases?

A. Amazon EC2, because EBS volumes can scale to hold any amount of data and, when used with Auto-scaling, can be designed for fault tolerance and high availability

B. Amazon S3, because it provides unlimited amounts of storage data, scales automatically, is highly available, and durable
C. Amazon Glacier, to keep costs low for storage and scale infinitely
D. Amazon Import/Export, because Amazon assists in migrating large amounts of data to Amazon S3

Answer: B

Explanation: The best option is to use S3 because you can host a large amount of data in S3 and is the best storage option provided by AWS.

https://docs.aws.amazon.com/AmazonS3/latest/dev/Welcome.html

361. When reviewing your auto-scaling events, you have noticed that the application is scaling up and down multiple times in the same hour, what design choice could you make to preserve elasticity with optimal cost? (Choose 2)

A. Modify the Auto-scaling group cool down timers

B. Modify the CloudWatch alarm period that triggers your Auto-scaling scale down policy
C. Modify the Auto-scaling group termination policy to terminate the newest instance first
D. Modify the Auto-scaling policy to use scheduled scaling actions

Answer: A and B

Explanation: The Auto-scaling cooldown period is a configurable setting for your Auto-scaling group that helps to ensure that Auto-scaling does not launch or terminate additional instances before the previous scaling activity takes effect. After the Auto-scaling group dynamically scales using a simple scaling policy, Auto-scaling waits for the cooldown period to complete before resuming scaling activities. When you manually scale your Auto-scaling group, the default is not to wait for the cooldown period, but you can override the default and honor the cooldown period. Note that, if an instance becomes unhealthy, Auto-scaling does not wait for the cooldown period to complete before replacing the unhealthy instance.

https://docs.aws.amazon.com/autoscaling/ec2/userguide/Cooldown.html

You can also modify the CloudWatch triggers to ensure the thresholds are appropriate for the scale down policy.

https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-scale-based-on-demand.html

362. Your startup company is going to launch the website in the

coming week. Chances are that the traffic will be very high in the beginning couple of weeks. If a load failure occurs, how can you set up DNS failover to a static website?

A. By adding more servers in case the application fails

B. By using Route 53 with failover option to failover to a static S3 website bucket or CloudFront distribution
C. By enabling failover to an on-premises data center to the app hosted there
D. By duplicating the exact application architecture in another region and configuring DNS weight-based routing

Answer: B

Explanation: Amazon Route53 health checks monitor the health and performance of your web applications, web servers, and other resources.

If you have multiple resources that perform the same function, you can configure DNS failover so that Amazon Route53 will route your traffic from an unhealthy resource to a healthy resource. For example, if you have two web servers and one web server becomes unhealthy, Amazon Route53 can route traffic to the other web server. So you can route traffic to a website hosted on S3 or to a CloudFront distribution.

https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover.html

363. For your organization need, you created a CloudFormation template to launch EC2 instances for both production and development environments in the same region. Each of the instance will have an Elastic IP and a security group associated with it. The development CloudFormation stack was created successfully but the production stack creation failed. What could be the reason?

A. You did not choose the Production version of the AMI you are using when creating the production stack

B. You hit the soft limit for security groups when creating the development environment
C. You hit the soft limit of 5 EIPs per region when creating the

development environment

D. You have chosen the wrong tags when creating the instances in both environments

Answer: C

Explanation: The most viable reason could be that you reached the limit for the number of Elastic IPs in the region. You can find more information about the EC2 service limits on the given URL:

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-resource-limits.html

364. For DR purposes, you have been assigned a task to build a duplicate environment in another region. A portion of your environment consists of EC2 instances with preconfigured software. Which of the below options will help you in configuring instances in another region?

A. Make the EC2 instance shareable among other regions through IAM permissions

B. Create an AMI of the EC2 instance

C. Create an AMI of the EC2 instance and copy the AMI to the desired region

D. None of the above

Answer: C

Explanation: You can copy an Amazon Machine Image (AMI) within or across an AWS region using the AWS Management Console, the AWS command line tools or SDKs, or the Amazon EC2 API, all of which support the CopyImage action. You can copy both Amazon EBS-backed AMIs and instance store-backed AMIs. You can copy AMIs with encrypted snapshots and encrypted AMIs.

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/CopyingAMIs.html

365. Your web application is running on six EC2 instances in an Auto-scaling group and it consumes about 45% of the resources on each instance. The number of requests to this application is consistent and does not experiences spikes. This is a mission-critical application

and you want high availability at all times. You want to evenly distribute the load between all instances. You also want to use the same AMI for all instances. Which of the following architectural choices should you make?

A. Deploy 2 EC2 instances in three regions and use Amazon Elastic Load Balancer

B. Deploy 3 EC2 instances in one availability zone and 3 in another availability zone and use Amazon Elastic Load Balancer
C. Deploy 3 EC2 instances in one region and 3 in another region and use Amazon Elastic Load Balancer
D. Deploy 6 EC2 instances in one availability zone and use Amazon Elastic Load Balancer

Answer: B

Explanation: For Option A and C, the ELB is designed to only run in one region in AWS and not across multiple regions. So these options are wrong.

Option D is automatically incorrect because remember that the question asks for high availability. For option A, if the AZ goes down then the entire application fails.

So for high availability and maintaining the number of instances, option B is correct. For more information on regions and AZs, visit the following URL

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.Regior

366. You have a set of web servers to host a web application on AWS. This app is used by a section of users and you want to monitor the number errors occurred when using the web application. Which of the following can be used to do this? (Choose 3)

A. Increase a metric filter in Cloudwatch whenever the pattern is matched

B. Search for the keyword “ERROR” in Cloudwatch logs
C. Search for the keyword “ERROR” in the log files on the server
D. Send the logs from the instances onto Cloudwatch logs

Answer: A, B, and D

Explanation: The AWS documentation mentions the following; You use metric filters to search for and match terms, phrases, or values in your log events. When a metric filter finds one of the terms, phrases, or values in your log events, you can increase the value of a CloudWatch metric. For example, you can create a metric filter to search for and count the occurrence of the word.

https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/FilterAndPatter

367. Your organization's web application is hosted on an Elastic BeanStalk environment and it is instructed that whenever application changes occur, and newer versions need to be deployed; the fastest deployment approach should be used. Which of the following deployment mechanisms fulfills this requirement?

A. Rolling with batch

B. Immutable
C. Rolling
D. All at once

Answer: D

Explanation: The requirement is to deploy the new version as fast as possible, according to AWS documentation, *All at Once* deployment mechanism is the fastest approach of deployment.

https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.deploy-existing-version.html

368. You are in charge of designing a CloudFormation template that deploys a LAMP stack. After deploying a stack, you notice that the Apache server is not up and running and is experiencing issues while starting up. However, the status of the stack is showing as CRETAE_COMPLETE. You want all resources defined in the stack to be up and running when the stack shows CREAT_COMPLETE status. How can you achieve this? (Choose 2)

A. By using the CFN helper scripts to signal once the resource configuration is complete

B. By using the CreationPolicy to ensure it is associated with the EC2 Instance resource
C. By using lifecycle hooks to mark the completion of the creation and configuration of the underlying resource
D. By defining a stack policy, which defines that all underlying resources should be up and running before showing a status of CREATE_COMPLETE

Answer: A and B

Explanation: When you provision an Amazon EC2 instance in an AWS CloudFormation stack, you might specify additional actions to configure the instance, such as install software packages or bootstrap applications. Normally, CloudFormation proceeds with stack creation after the instance has been successfully created. However, you can use a CreationPolicy so that CloudFormation proceeds with stack creation only after your configuration actions are done. That way you will know your applications are ready to go after stack creation succeeds.

https://aws.amazon.com/blogs/devops/use-a-creationpolicy-to-wait-for-on-instance-configurations/

369. Your company's application is hosted on AWS on EC2 instances. The IT security department has given a requirement to process and analyze the logs of these instances in real time. Which of the following can be used for this?

A. Another EC2 Instance with a larger instance type to process the logs
B. Amazon S3 to store the logs and then Amazon Kinesis to process and analyze the logs in real time
C. Amazon Glacier to store the logs and then Amazon Kinesis to process and analyze the logs in real time
D. Cloudwatch logs to process and analyze the logs in real time

Answer: B

Explanation: The AWS documentation says:

Real-time Metrics and Reporting

You can use data collected into Kinesis Streams for simple data analysis and

reporting in real time. For example, your data-processing application can work on metrics and reporting for system and application logs as the data is streaming in, rather than waiting to receive batches of data.

Real-time Data Analytics

This combines the power of parallel processing with the value of real-time data. For example, process website clickstreams in real time, and then analyze site usability engagement using multiple different Kinesis Streams applications running in parallel.

Amazon Glacier is meant for Archival purposes and should not be used for storing the logs for real time processing.

https://docs.aws.amazon.com/streams/latest/dev/introduction.html

370. Which of the following deployment types are available in AWS CodeDeploy? (Choose 2)

A. Blue/Green Deployments
B. Immutable Deployments
C. Rolling Deployments
D. In-place Deployments

Answer: A and D

Explanation: The AWS documentation mentions the following:

Deployment Type: The method used to make the latest application revision available on instances in a deployment group.

- *In-place Deployment:* The application on each instance in the deployment group is stopped, the latest application revision is installed, and the new version of the application is started and validated. You can choose to use a load balancer so each instance is deregistered during its deployment and then restored to service after the deployment is complete
- *Blue/Green Deployment:* The instances in a deployment group (the original environment) are replaced by a different set of instances (the replacement environment) using these steps:
    - Instances are provisioned for the replacement environment
    - The latest application revision is installed on the replacement

instances
    - An optional wait time occurs for activities such as application testing and system verification
    - Instances in the replacement environment are registered with an Elastic Load Balancing load balancer, causing traffic to be rerouted to them. Instances in the original environment are deregistered and can be terminated or kept running for other uses

https://docs.aws.amazon.com/codedeploy/latest/userguide/primary-components.html

371. Your company has a number of applications that need to be migrated to AWS. Initially you thought to move these to Elastic BeanStalk service, but you noticed that the underlying platform service is not an option in the Elastic BeanStalk environment. Which of the following options can be used to move your applications to Elastic BeanStalk?

A. Create a Docker container for the custom application and then deploy it to Elastic BeanStalk
B. Use custom CloudFormation templates to deploy the application into Elastic BeanStalk
C. Use custom Chef recipes to deploy your application in Elastic BeanStalk
D. Use the OpsWorks service to create a stack. In the stack, create a separate custom layer. Deploy the application to this layer and then attach the layer to Elastic BeanStalk

Answer: A

Explanation: Elastic Beanstalk supports the deployment of web applications from Docker containers. With Docker containers, you can define your own runtime environment. You can choose your own platform, programming language, and any application dependencies (such as package managers or tools), that are not supported by other platforms. Docker containers are self-contained and include all the configuration information and software your web application requires to run.

https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker.

372. Your web application is hosted on EC2 instances in an Auto-scaling group that are sitting behind an ELB. You have created AMIs of each application version for deployment and you want to deploy the newer version of your application using Blue/Green deployment technique. You have chosen the Blue/Green deployment model because you want to migrate users in a controlled manner while the size of the fleet remains constant for at least six hours to ensure that there are no issues in the new version. What would you do to enable this technique while being able to roll back easily? (Choose 2)

A. Configure Elastic Load Balancing to vary the proportion of requests sent to instances running the two application versions
B. Use Amazon Route53 weighted Round Robin to vary the proportion of requests sent to the load balancers
C. Create an Auto-scaling launch configuration with the new AMI to use the new launch configuration and to register instances with the existing load balancer
D. Create an Auto-scaling launch configuration with the new AMI to use the new launch configuration and to register instances with the new load balancer

Answer: B and D

Explanation: The AWS documentation describes Blue/Green deployment as the following; You can shift traffic all at once or you can do a weighted distribution. With Amazon Route 53, you can define a percentage of traffic to go to the green environment and gradually update the weights until the green environment carries the full production traffic. A weighted distribution provides the ability to perform canary analysis where a small percentage of production traffic is introduced to a new environment. You can test the new code and monitor for errors, limiting the blast radius if any issues are encountered. It also allows the green environment to scale out to support the full production load if you are using Elastic Load Balancing.
https://d0.awsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf

373. You are a DevOps Engineer in an organization whose

application is hosted on a single EC2 instance in AWS. You are receiving complaints from end users about slow response time of your application. What would you do to resolve this issue?

A. By using CloudFormation to deploy the app again with an Amazon RDS with the Multi-AZ feature
B. By using Amazon RDS with the Multi-AZ feature
C. By using Auto-scaling launch configurations to launch multiple instances and placing them behing an ELB
D. By using Auto-scaling groups to launch multiple instances and placing them behind an ELB

Answer: A

Explanation: When you use Auto-scaling, you can automatically increase the size of your Auto-scaling group when demand goes up and decrease it when demand goes down. As Auto-scaling adds and removes EC2 instances, you must ensure that the traffic for your application is distributed across all of your EC2 instances. The Elastic Load Balancing service automatically routes incoming web traffic across such a dynamically changing number of EC2 instances. Your load balancer acts as a single point of contact for all incoming traffic to the instances in your Auto-scaling group.
https://docs.aws.amazon.com/autoscaling/ec2/userguide/autoscaling-load-balancer.html

374. You are designing a CloudFormation template to launch EC2 instances as web servers and following user data needs to be passed to the instances:

```
#!/bin/bash
sudo apt-get update
sudo apt-get install -y nginx
```

In which portion of the template should you pass this data?

A. In the Metadata section of the EC2 Instance in the Output section
B. In the Metadata section of the EC2 Instance in the resources section
C. In the properties section of the EC2 Instance in the Output

section

D. In the properties section of the EC2 Instance in the resources section

Answer: D

Explanation: The following is an example of how user data can be passed to instances using a CloudFormation template:

```
{
    "Resources": {
        "WebServerInstance": {
        "Type": "AWS::EC2::Instance".
        "Properties": {
        "InstanceType": "t2.micro".
        "ImageId": "ami-6f198a0c".
        "UserData": {
        "Fn::Base64":.

        "Fn::Join": {
        "\n".
        [
        "#!/bin/bash".
        "sudo apt-get update".
        "sudo apt-get install –y nginx]]}}
}
}
}
}
```

https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/deployi

375. While using the Elastic Load Balancer, which of the following are ways to secure data in-transit? (Choose 2)

A. Using an HTTPS front end listener for your ELB
B. Using an HTTP front end listener for your ELB
C. Using an SSL front end listener for your ELB
D. Using a TCP front end listener for your ELB

Answer: A and C

Explanation: As per AWS documentation; You can create a load balancer that uses the SSL/TLS protocol for encrypted connections (also known as SSL offload). This feature enables traffic encryption between your load balancer and the clients that initiate HTTPS sessions, and for connections between your load balancer and your EC2 instances.

If you use HTTPS or SSL for your front-end connections, you must deploy an X.509 certificate (SSL server certificate) on your load balancer. The load balancer decrypts requests from clients before sending them to the back-end instances (known as SSL Termination).

If you do not want the load balancer to handle the SSL termination (known as SSL offloading) you can use TCP for both the front-end and back-end connections, and deploy certificates on the registered instances handling requests.

https://docs.aws.amazon.com/elasticloadbalancing/latest/classic/elb-https-load-balancers.html

https://docs.aws.amazon.com/elasticloadbalancing/latest/classic/elb-listener-config.html

https://docs.aws.amazon.com/elasticloadbalancing/latest/classic/using-elb-listenerconfig-quickref.html

376. To detect the application health while performing a Blue/Green deployment, which of the following services can be used?

A. AWS CloudTrail
B. AWS CloudWatch
C. AWS CodeStar
D. AWS CodeCommit

Answer: B

Explanation: Amazon CloudWatch is a monitoring service for AWS Cloud resources and the applications you run on AWS. CloudWatch can collect and track metrics, collect and monitor log files, and set alarms. It provides system-wide visibility into resource utilization, application performance, and operational health, which are key to early detection of application health in Blue/Green deployments.

https://d0.awsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf

377. Your company's web application is hosted on a set of EC2 instances sitting behind an ELB. An Amazon RDS instance is also used by this app. You are required to make this infrastructure self-healing and cost effective. Which of the following would fulfil this requirement? (Choose 2)

A. Use CloudWatch metrics to check the utilization of the databases servers. Use Auto-scaling group to scale the database instances accordingly based on the CloudWatch metrics
B. Utilize the Read Replica feature for the Amazon RDS layer
C. Use CloudWatch metrics to check the utilization of the web layer. Use Auto-scaling group to scale the web instances accordingly based on the CloudWatch metrics
D. Utilize the Multi-AZ feature for the Amazon RDS layer

Answer: C and **D**

Explanation: Self-healing architecture for the scenario in question can be achieved by using Auto-scale.

Amazon RDS Multi-AZ deployments provide enhanced availability and durability for Database (DB) Instances, making them a natural fit for production database workloads. When you provision a Multi-AZ DB Instance, Amazon RDS automatically creates a primary DB Instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). Each AZ runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. In case of an infrastructure failure, Amazon RDS performs an automatic failover to the standby (or to a read replica in the case of Amazon Aurora), so that you can resume database operations as soon as the failover is complete. Since the endpoint for your DB Instance remains the same after a failover, your application can resume database operation without the need for manual administrative intervention.

https://aws.amazon.com/rds/details/multi-az/

378. A group of users in your organization uses a set of instances that is hosting nginx server and a web application. The instances start facing technical issues after a recent version upgrade and require

immediate restart. You did not get time to inspect what caused the issue on the servers. Which of the following options, if implemented prior to the incident would have assisted in you in finding out the issue?

A. Streaming all the data to Amazon Kinesis and then analyzing the data in real time
B. Enabling detailed monitoring and checking the Cloudwatch metrics to see the cause of the issue
C. Installing Cloudwatch logs agent on the instance and sending all the logs to Cloudwatch logs
D. Creating a snapshot of the EBS volume before restart, attaching it to another instance as a volume and then diagnosing the issue

Answer: C

Explanation: You can publish log data from Amazon EC2 instances running Linux or Windows Server, and logged events from AWS CloudTrail. CloudWatch Logs can consume logs from resources in any region, but you can only view the log data in the CloudWatch console in the regions where CloudWatch Logs is supported.

Option A is incorrect as here we are dealing with an issue concerning the underlying application that handles the data and so this solution will not help.

Option B is invalid as detailed monitoring will only help us to get more information about the performance metrics of the instances, volumes etc. and will not be able to provide full information regarding technical issues.

Option D is incorrect as if we had created a snapshot prior to the update, it might have been useful but not after the incident.

https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/StartTheCWLA

379. You are required to deploy a multi-container Docker environment to Elastic BeanStalk. Which of the following files can be used to deploy a set of Docker containers in BeanStalk?

A. Dockerrun
B. Dockerrun.aws.json
C. DocekrMultiFile
D. DockerFile

Answer: B

Explanation: The AWS documentation states; A Dockerrun.aws.json file is an Elastic BeanStalk–specific JSON file that describes how to deploy a set of Docker containers as an Elastic BeanStalk application. You can use a Dockerrun.aws.json file for a multi-container Docker environment.

Dockerrun.aws.json describes the containers to deploy to each container instance in the environment as well as the data volumes to create on the host instance for the containers to mount.

https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker

380. Which of the following can be integrated with Jenkins continuous integration tool?

A. Amazon EC2
B. Amazon ECS
C. Amazon Elastic BeanStalk
D. All of the above

Answer: D

Explanation: The following AWS services can be integrated with Jenkins:

- Amazon EC2
- Amazon ECR
- Amazon SNS
- Amazon ECS
- Amazon S3
- AWS CloudFormation
- AWS CodeDeploy
- AWS Code PipeLine
- AWS CodeCommit
- AWS Device Farm
- AWS Elastic BeanStalk

381. Your company's application is hosted on AWS and uses DynamoDB. According to the IT security policy, it is required to record all the source IP addresses that make calls to DynamoDB table. Which of the following service can be used to fulfil this requirement?

A. AWS CloudWatch
B. AWS Trusted Advisor
C. AWS CloudTrail
D. AWS Inspector

Answer: C

Explanation: DynamoDB is integrated with CloudTrail, a service that captures low-level API requests made by or on behalf of DynamoDB in your AWS account and delivers the log files to an Amazon S3 bucket that you specify. CloudTrail captures calls made from the DynamoDB console or from the DynamoDB low-level API. Using the information collected by CloudTrail, you can determine what request was made to DynamoDB, the source IP address from which the request was made, who made the request, when it was made, and so on.

https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/logging using-cloudtrail.html

382. Your organization's infrastructure is built in a way that the instances access data objects that are stored in an S3 bucket. The IT security department is concerned about the security of the architecture and you are required to implement the following:

- Ensure that the EC2 instance securely access data objects in the S3 bucket
- Ensure that the integrity of the objects stored in S3 is maintained

Which of the following would help you in fulfilling these requirements? (Choose 2)

A. Use S3 Cross Region replication to replicate the objects so that the integrity of data is maintained
B. Create an IAM Role and ensure the EC2 Instance uses the IAM Role to access the data in the bucket
C. Use an S3 bucket policy that ensures that MFA Delete is set on the objects in the bucket
D. Create an IAM user and ensure the EC2 Instances uses the IAM user credentials to access the data in the bucket

Answer: B and C

Explanation: IAM roles are designed so that your applications can securely make API requests from your instances, without requiring you to manage the security credentials that the applications use. Instead of creating and distributing your AWS credentials, you can delegate permission to make API requests using IAM roles.

MFS Delete can be used to add another layer of security to S3 Objects to prevent accidental deletion of objects.

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html

https://aws.amazon.com/blogs/security/securing-access-to-aws-using-mfa-part-3/

383. Which of the below mentioned source repositories can be used by AWS CodeDeploy to deploy code? (Choose 3)

A. S3 buckets
B. GitHub Repositories
C. Subversion Repositories
D. Bitbucket Repositories

Answer: A, C, and D

Explanation: You can deploy a nearly unlimited variety of application content, such as code, web and configuration files, executables, packages, scripts, multimedia files, and so on. AWS CodeDeploy can deploy application content stored in Amazon S3 buckets, GitHub repositories, or Bitbucket repositories. You do not need to make changes to your existing code before you can use AWS CodeDeploy.

https://docs.aws.amazon.com/codedeploy/latest/userguide/welcome.html

384. You want to get a snapshot of the current configuration of the resources in your company's AWS account. Which of the following can be used for this?

A. AWS IAM
B. AWS Trusted Advisor
C. AWS Config
D. AWS Trusted Advisor

Answer: C

Explanation: With AWS Config, you can do the following:

- Evaluate your AWS resource configurations for desired settings
- Get a snapshot of the current configurations of the supported resources that are associated with your AWS account
- Retrieve configurations of one or more resources that exist in your account
- Retrieve historical configurations of one or more resources
- Receive a notification whenever a resource is created, modified, or deleted
- View relationships between resources. For example, you might want to find all resources that use a particular security group

https://docs.aws.amazon.com/config/latest/developerguide/WhatIsConfig.htm