Answer Sheet

Q1) Robert wants to have a snapshot of the current configuration of the resources in the AWS account. Select the AWS service to accomplish it.

- None of these
- AWS Code Deploy
- AWS Trusted Advisor
- AWS LAM
- ✓ AWS Config

Q2) Which of the following Deployment types are available in the Code Deploy service. Choose 2 answers from the options given below

- ✓ Blue/green deployment
- Immutable deployment
- Rolling deployment
- ✓ In-place deployment

Q3)

You work for a very large company that has multiple applications which are very different and built on different programming languages.

How can you deploy applications as quickly as possible?

- Develop each app in a separate Docker containers and deploy using Cloud Formation
- ✓ Develop each app In a separate Docker container and deploy using Elastic Beanstalk
- Create a Lambda function deployment package consisting of code and any dependencies
- Develop each app in one Docker container and deploy using Elastic Beanstalk

---

Q4)

You are the IT administrator for your company. You have the responsibility of creating development environments which would confirm to the LAMP development stack. The requirement is that the development team always gets the latest version of the LAMP stack each time a new instance is launched.

Which of the following is an efficient and effective way to implement this requirement? Choose 2 answers from the options given below

- Create an AMI with all the artifacts of the LAMP stack and provide an instance to the development team based on the AMI.)
- ✓ Create a cloud formation template and use the cloud-mit directives to download and the install the LAMP stack packages.
- ✓ Use the User data section and use a custom script which will be used to download the necessary LAMP stack packages.
- Create an EBS Volume with the LAMP stack and attach it to an instance whenever it Is required.

Q5)

You need your CI to build AMIs with code pre-installed on the images on every new code push. You need to do this as cheaply as possible.

How do you do this?

- 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 ondemand instance.
- When the CI Instance receives commits, attach a new EBS volume to the CI machine. Perform all setup on this EBS volume so you don't need a new EC2 instance to create the AMP.
- ✓ Bid on spot instances just above the asking price as soon as new commits come in. perform all instance configuration and setup, then create an AMP based on the spot instance.
- Purchase a Light Utilization Reserved Instance to save money on the continuous integration machine. Use these credits whenever your create AMPS Ofl instances.

Q6) When creating an Elastic Beanstalk environment using the Wizard , what are the 3 configuration options presented to you ?

- ✓ Choosing the type of Notification SNS or SQS
- Choosing whether you want a highly available environment or not
- ✓ Choosing the type of Environment - Web or Worker environment
- ✓ Choosing the platform type - Node.js. 115, etc

Q7) Which of the below services can be used to deploy application code content stored in Amazon S3 buckets, Git Hub repositories, or Bit bucket repositories ?

- Route53
- Code Commit
- ✓ Code Deploy
- 53 Lifecycles

Q8)

Your company has a set of resources hosted in AWS. They want to be notified when the costs of the AWS resources running in the account reaches a certain threshold.

How can this be accomplished in an ideal way?

- Create a script which monitors all the running resources and calculates the costs accordingly.
- Download the cost reports and analyze the reports to see if the costs are going beyond the threshold
- Create a consolidated biting report and see if the costs are going beyond the threshold.
- ✅ Create a billing alarm which can alert you when the costs are going beyond a certain threshold ?

Q9)

You currently have an Auto scaling group that has the following settings Mm capacity- 2 Desired capacity - 2 Maximum capacity -4 The current number of instances running in the Auto scaling Group is 2.

You have been notified that for a duration of an hour, you need to ensure that no new Instances are launched by the Auto scaling Group Which of the below 2 actions can be carried out to fulfill this requirement

- ✅ Suspend the Launch process of the Auto scaling Group
- Change the Minimum capacity to 2
- Change the Desired capacity to 4
- ✅ Change the Maximum capacity to 2

Q10) Which of the following credentials types are supported by AWS Code Commit?

- ✅ SSH keys
- ✅ Username/password
- ✅ Git Credentials
- AWS Access keys

Q11) Which of the following design strategies is ideal when designing loosely coupled systems. Choose 2 answers from the options given below ?

- Having the web and worker roles running on the same set of EC2 Instances
- Using SNS to establish communication between the web and worker roles
- ✅ Having the web and worker roles running on separate EC2 Instances
- ✅ Using SQS to establish communication between the web and worker roles

---

Q12)

You are in charge of creating a Cloud formation template that will be used to spin our resources on demand for your Devops team. The requirement is that this cloud formation template should be able to spin up resources In different regions.

Which of the following aspects of Cloud formation templates can help you design the template to spin up resources based on the region?

- Use the metadata section in the Cloud formation template. so that based on the relevant region, the relevant resource can be spinned up.
- Use the outputs section in the Cloud formation templates so that based on the relevant region , the relevant resource can be spinned up.
- Use the parameters section in the Cloud formation template. so that based on the relevant region, the relevant resource can be spinned up.
- ✅ Use mappings section in the Cloud formation template . so that based on the relevant region, the relevant resource can be spinned up.

Q13)

You have launched a cloud formation template, but are receiving a failure notification after the template was launched.

What is the default behavior of Cloud formation in such a case Please select:

- It will continue with the creation of the next resource in the stack
- It will keep all the resources that were created up to the failure point.
- It will prompt the user on whether to keep or terminate the already created resources
- ✅ It will rollback all the resources that were created up to the failure point

---

Q14)

An enterprise wants to use a third-party SaaS application running on AWS.. The SaaS application needs to have access to issue several API commands to discover Amazon EC2 resources running within the enterprise?s account. The enterprise has internal security policies that require any outside access to their environment must conform to the principles of least privilege and there must be controls in place to ensure that the credentials used by the SaaS vendor cannot be used by any other third party.

Which of the following would meet all of these conditions?

- Create an IAM role for EC2 instances, assign it a policy that allows only the actions required tor the Saas application to work, provide the role ARN to the SaaS provider to use when launching their application instances.
- From the AWS Management Console, navigate to the Security Credentials page and retrieve the access an secret key for your account.
- Create an IAM user within the enterprise account assign a user policy to the IAM user that allows only the actions required by the SaaS application. Create a new access and secret key for the user and provide these credentials to the SaaS provider.
- ✅ Create an IAM role for cross-account access allows the SaaS providers account to assume the role and assign It a policy that allows only the actions required by the SaaS application.

Q15)

You are working as an AWS Devops admins for your company. You are in-charge of building the Infrastructure for the company's development teams using Cloud formation. The template will include building the VPC and networking components, installing a LAMP stack and securing the created resources.

As per the AWS best practices what is the best way to design this template ?

- Create a single cloud formation template to create all the resources since it would be easier from the maintenance perspective.
- Create multiple cloud formation templates based on the number of development groups In the environment.
- ✓ Create multiple cloud formation templates for each set of logical resources, one for networking, tne otner for LAMP stack creation.
- Create multiple cloud formation templates based on the number of VPCs in the environment.

Q16) When building a multi container Docker platform using Elastic Beanstalk , which of the following is required Please select:

- Kurbernetes to manage the docker containers.
- ✓ Prebuilt Images stored in a public or private online image repository.
- Red Hat Opensift to manage the docker containers.
- Docker File to create custom images during deployment

Q17)

You currently have EC2 Instances hosting an application. These instances are part of an Autoscaling Group. You now want to change the instance type of the EC2 Instances.

How can you manage the deployment with the least amount of downtime ?

- Use the Rolling Update feature which is available for EC2 Instances.
- ✓ Use the Auto Scaling Rolling Update policy on Cloud Formation Template Auto Scaling group
- Manually terminate the Instances, launch new Instances with the new Instance type and attach them to the
- Terminate the existing Auto Scaling group. Create a new launch configuration with the new Instance type. Attach that to the new Auto scaling Group.

Q18) When your application is loaded onto an Ops works stack, which of the following event is triggered by Show Ops works?

- ✓ Deploy Fill
- Configure
- Shutdown
- Setup

Q19)

You have carried out a deployment using Elastic Beanstalk, but the application is unavailable.

What could be the reason for this ?

- The cool down period is not properly configured for Elastic Beanstalk
- You need to configure ELB along with Elastic Beanstalk
- ✓ There will always be a few seconds of downtime before the application is available
- You need to configure Route53 along with Elastic Beanstalk

Q20)

Your company has a set of EC2 resources hosted on AWS. Your new IT procedures state that AWS EC2 Instances must be of a particular Instance type.

Which of the following can be used to get the list of EC2 Instances which currently don't match the instance type specified in the new IT procedures

- Use VPC Flow Logs to check which EC2 Instances don't match the intended instance type.
- Use Trusted Advisor to check which EC2 Instances don't match the intended instance type.
- ✓ Use AWS Config to create a rule to check the EC2 Instance type
- Use AWS Cloud watch alarms to check which EC2 Instances don't match the intended instance type.

Q21)

You are currently using Elastic Beanstalk to host your production environment. You need to rollout updates to your application hosted on this environment. This is a critical application which is why there is a requirement that the rollback, if required . should be carried out with the least amount of downtime.

Which of the following deployment strategies would ideally help achieve this purpose Please select:

- Create a Cloud formation template with the same resources as those in the Elastic beanstalk environment, if the deployment fails, deploy the Cloud formation template.
- Create another parallel environment in elastic beanstalk. Create a new Route53 Domain name for the new environment and release that un to the users.
- Use Rolling updates in Elastic Beanstalk so that if the deployment fails , the rolling updates feature would roll back to the last deployment.
- ✓ Create another parallel environment in elastic beanstalk. Use the Swap URL feature.

Q22)

Your company has a set of resources hosted in AWS. Your IT Supervisor is concerned with the costs being incurred by the resources running in AWS and wants to optimize on the costs as much as possible.

Which of the following ways could help achieve this efficiently? Choose 2 answers from the options given below?

- Create Cloud watch logs to monitor underutilized resources and either shutdown or terminate resources which are not required.
- ✓ Create Cloud watch alarms to monitor underutilized resources and either shutdown or terminate resources which are not required.
- ✓ Use the Trusted Advisor to see underutilized resources
- Create a script which monitors all the running resources and calculates the costs accordingly. The analyze those resources accordingly and see which can be optimized.

Q23)

You are deciding on a deployment mechanism for your application.

Which of the following deployment mechanisms provides the fastest rollback after failure?

- Rolling- Immutable
- ✓ Blue/Green
- Rolling. Mutable
- Canary

Q24)

An audit is going to be conducted for your company's AWS account.

Which of the following steps will ensure that the auditor has the right access to the logs of your AWS account ?

- Enable Cloud watch logs. Create a user for the IT Auditor and ensure that full control Is given to the user for the Cloud watch logs.
- ✓ Ensure that Cloud trail is enabled. Create a role for read only access to Cloud trail. Create a user for the IT Auditor and attach the role to the user.
- Enable 53 and ELB logs. Send the logs as a zip file to the IT Auditor
- Ensure that Cloud trail is enabled. Create a user for the IT Auditor and ensure that full control is given to th user for Cloud trail.

Q25) Which of the following environment types are available in the Elastic Beanstalk environment. Choose 2 answers from the options given below?

- ✓ Single Instance
- Multi-Instance
- ✓ Load Balancing, Auto scaling
- SQS. Auto scaling

Q26)

Your company is planning to setup a word press application. The word press application will connect to a My SQL database.

Part of the requirement is to ensure that the database environment is fault tolerant and highly available. Which of the following 2 options individually can help fulfill this requirement?

- Create multiple [C2 instances in the same AZ. Host MYSQL and enable replication via scripts between the instances.
- Create a My SQL RDS environment and create a Read Replica
- ✓ Create multiple EC2 Instances In separate Ars. Host MySQL and enable replication via scripts between the instances.
- ✓ Create a MYSQL RDS environment with Multi-AZ feature enabled

Q27)

You have a code repository that uses Amazon 53 as a data store. During a recent audit of your security controls, some concerns were raised about maintaining the integrity of the data in the Amazon S3 bucket. Another concern was raised around securely deploying code from Amazon S3 to applications running on Amazon EC2 in a virtual private cloud.

What are some measures that you can implement to mitigate these concerns? Choose two answers from the options given below?

- ✓ Create an Amazon Identity and Access Management role with authorization to access the Amazon 53 bucket, and launch all of your application's Amazon EC2 instances with this role.
- Use AWS Data Pipeline with multi-factor authentication to securely deploy code from the Amazon S3 bucket to your Amazon EC2 instances.
- 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.
- ✓ Add an Amazon S3 bucket policy with a condition statement that requires multi-factor authentication in order to delete objects and enable bucket version ing.

Q28)

You have an application running on Amazon EC2 in an Auto Scaling group. Instances are being bootstrapped dynamically, and the bootstrapping takes over 15 minutes to complete. You find that instances are reported by Auto Scaling as being In Service before bootstrapping has completed. You are receiving application alarms related to new instances before they have completed bootstrapping, which is causing 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?

- Use the default Amazon Cloud Watch application metrics to monitor your application's health. Configure an Amazon SNS topic to send these Cloud Watch alarms to the correct recipients.
- ✓ 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.
- Increase the desired number of instances in your Auto Scaling group configuration to reduce the time it takes to bootstrap future instances.
- 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.

Q29)

When an Auto Scaling group is running in Amazon Elastic Compute Cloud (EC2), your application rapidly scales up and down in response to load within a 1 0-minute window; however, after the load peaks, you begin to see problems in your configuration management system where previously terminated Amazon EC2 resources are still showing as active.

What would be a reliable and efficient way to handle the cleanup of Amazon EC2 resources within your configuration management system? Choose two answers from the options given below

- ✓ Write a script that is run by a daily cron job on an Amazon EC2 instance and that executes API Describe c of the EC2 Auto Scaling group and removes terminated instances from the configuration management system.
- Use your existing configuration management system to control the launching and bootstrapping of instances to reduce the number of moving parts in the automation.
- ✓ Write a small script that is run during Amazon EC2 instance shutdown to de-register the resource from the configuration management system.
- 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.

Q30)

You are a Dev Ops engineer for a company. You have been requested to create a rolling deployment solution that is cost-effective with minimal downtime.

How should you achieve this? Choose two answers from the options below

- After each stack is deployed, tear down the old stack
- Re-deploy your application using a Cloud Formation template to deploy Elastic Beanstalk
- ✓ Use Update Policy attribute to specify how Cloud Formation handles updates to Auto Scaling Group resource
- ✓ Re-deploy with a Cloud Formation template. define update policies on Auto Scaling groups in your Cloud Formation template

---

Q31)

You have a complex system that involves networking, IAM policies, and multiple, threetier applications. You are still receiving requirements for the new system, so you don't yet know how many AWS components will be present in the final design. You want to start using AWS Cloud Formation to define these AWS resources so that you can automate and version-control your infrastructure.

How would you use AWS Cloud Formation to provide agile new environments for your customers in a cost-effective, reliable manner?

- ✓ Create multiple separate templates for each logical part of the system, create nested stacks in AWS Cloud Formation. and maintain several templates to version-control.
- Manually construct the networking layer using Amazon Virtual Private Cloud (VPC) because this does not change often, and then use AWS Cloud Formation to define all other ephemeral resources.
- Manually create one template to encompass all the resources that you need for the system, only have a single template to version-control.
- Create multiple separate templates for each logical part of the system. and provide the outputs from one to the next using an Amazon Elastic Compute Cloud (EC2) Instance running the SDK for finer granularity of control.

Q32)

You have an ELB setup in AWS with EC2 instances running behind them.

You have been requested to monitor the incoming connections to the ELB. Which of the below options can suffice this requirement?

- ✓ Create a custom metric Cloud Watch filter on your load balancer
- Use AWS Cloud Trail with your load balancer
- Enable access logs on the load balancer
- Use a Cloud Watch Logs Agent

Q33)

You are doing a load testing exercise on your application hosted on AWS. While testing your Amazon RDS My SQL DB Instance, you notice that when you hit 100% CPU utilization on it. your application becomes non- responsive. Your application Is read-heavy.

What are methods to scale your data tier to meet the application's needs? Choose three answers from the options given below Please select:

✅ Use Elastic Cache In front of your Amazon RDS DB to cache common queries.
✅ Add Amazon RDS DB read replicas, and have your application direct read queries to them.
⚪ Use an Amazon SQS queue to throttle data going to the Amazon RDS DB instance.
✅ Shard your data set among multiple Amazon RDS DB instances.

Q34)

Your company develops a variety of web applications using many platforms and programming languages with different application dependencies. Each application must be developed and deployed quickly and be highly available to satisfy your business requirements.

Which of the following methods should you use to deploy these applications rapidly?

⚪ Develop each application's code In Dynamo DB. and then use hooks to deploy it to Elastic Beanstalk environments with Auto Scaling and Elastic Load Balancing.
✅ Develop the applications In Ducker containers, and then deploy them to Elastic Beanstalk environments with Auto Scaling and Elastic Load Balancing.
⚪ Store each application's code in a Git repository, develop custom package repository managers for each application's dependencies. and deploy to AWS Ops Works in multiple Availability Zones.
⚪ Use the AWS Cloud Formation Ducker import service to build and deploy the applications with high availability in multiple Availability Zones.

Q35)

You have an application consisting of a stateless web server tier running on Amazon EC2 instances behind load balancer, and are using Amazon RDS with read replicas.

Which of the following methods should you use to implement a self-healing and cost-effective architecture? Choose 2 answers from the options given below

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

Q36)

You are administering a continuous integration application that polls version control for changes and then launches new Amazon EC2 instances for a full suite of build tests.

What should you do to ensure the lowest overall cost while being able to run as many tests in parallel as possible?

⚪ Perform all tests on the continuous integration system. using AWS Ops Works for unit. integration, and build tests.
⚪ Perform syntax checking on the continuous integration system before launching a new Amazon EC2 instance for build test.. unit and Integration tests.
✅ Perform syntax and build tests on the continuous integration system before launching the new Amazon E instance unit and integration tests.
⚪ 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.

Q37) You have a multi-docker environment that you want to deploy to AWS. Which of the following configuration files can be used to deploy a set of Docker containers as an Elastic Beanstalk application? Please select:

⚪ .extensions
✅ Ducker run .aws. json
⚪ Ducker run.json
⚪ Ducker file

Q38)

One of the instances in your Auto Scaling group health check returns the status of Impaired to Auto Scaling.

What will Auto Scaling do in this case?

✅ Terminate the instance and launch a new instance
⚪ Perform a health check until cool down before declaring that the instance has failed
⚪ Send an SNS notification
⚪ Wait for the instance to become healthy before sending traffic

---

Q39)

You have just recently deployed an application on EC2 instances behind an ELB. After a couple of weeks, customers are complaining on receiving errors from the application. You want to diagnose the errors and are trying to get errors from the ELB access logs. But the ELB access logs are empty.

What Is the reason for this. Please select:

✅ Access logging is an optional feature of Elastic Load Balancing that is disabled by default .
⚪ ELB Access logs are only available for a maximum of one week.

- You do not have the appropriate permissions to access the logs
- You do not have your Cloud Watch metrics correctly configured

---

Q40)

You are using a configuration management system to manage your Amazon EC2 instances. On your Amazon EC2 Instances, you want to store credentials for connecting to an Amazon RDS DB instance.

How should you securely store these credentials?

- Store the Amazon RDS DB credentials in Amazon EC2 user data. Import the credentials Into the Instance on boot.
- Give the Amazon EC2 instances an PAM role that allows read access to a private Amazon S3 bucket. Store a file with database credentials In the Amazon 53 bucket. Have your configuration management system pull the file from the bucket when it is needed.
- ✅ 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.
- Launch an Amazon EC2 instance and use the configuration management system to bootstrap the instanc with the Amazon RDS DB credentials. Create an AMP from this instance.

Q41) You have the following application to be setup in AWS 1) A web tier hosted on EC2 Instances 2) Session data to be written to Dynamo DB 3) Log files to be written to Microsoft SQL Server How can you allow an application to write data to a Dynamo DB table?

- Add an IAM user to a running EC2 instance.
- ✅ Create an IAM role that allows write access to the Dynamo DB table.
- Create an IAM role that allows read access to the Dynamo DB table.
- Add an AM user that allows write access to the Dynamo DB table.

Q42)

Your mobile application includes a photo-sharing service that is expecting tens of thousands of users at launch. You will leverage Amazon Simple Storage Service (S3) for storage of the user Images, and you must decide how to authenticate and authorize your users for access to these images. You also need to manage the storage of these images.

Which two of the following approaches should you use? Choose two answers from the options below

- Authenticate your 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 users objects to that bucket.
- ✅ Use a key-based naming scheme comprised from the user IDs for all user objects in a single Amazon S3 bucket.
- Authenticate your users at the application level, and use AWS Security Token Service (5Th) to grant token- based authorization to S3 objects.
- Use AWS Identity and ACCeSS Management (IAM) user accounts as your applicationlevel user database, and offload the burden of authentication from your application code.

Q43)

You currently have an Auto Scaling group with an Elastic Load Balancer and need to phase out all instances and replace with a new instance type.

What are 2 ways in which this can be achieved. Please select:

- Attach an additional ELB to your Auto Scaling configuration and phase in newer instances while removing older Instances
- ✅ Use News instance to phase out all instances that use the previous configuration.
- ✅ Attach an additional Auto Scaling configuration behind the ELB and phase in newer instances while removing older instances. ..
- Use Oldest Launch Configuration to phase out all instances that use the previous configuration.

Q44)

You have an Auto Scaling group of Instances that processes messages from an Amazon Simple Queue Service (SQS) queue. The group scales on the size of the queue. Processing Involves calling a third..party web service. The web service Is complaining about the number of failed and repeated calls it is receiving from you. You have noticed that when the group scales in, instances are being terminated while they are processing.

What cost-effective solution can you use to reduce the number of incomplete process attempts?

- ✅ 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.
- Increase the minimum and maximum size for the Auto Scaling group. and change the scaling policies so they scale less dynamically.
- Create a new Auto Scaling group with minimum and maximum of 2 and instances running web proxy software. Configure the VPC route table to route HTTP traffic to these web proxies.
- Modify the application running on the instances to enable termination protection while it processes a task and disable it when the processing is complete.

Q45)

You are using Elastic Beanstalk to manage your e-commerce store. The store is based on an open source ec commerce platform and is deployed across multiple instances In an Auto Scaling group. Your development team often creates new ? extensions? for the e-commerce store. These extensions include PHP source code as well as an SQL upgrade script used to make any necessary updates to the database schema. You have noticed that some extension deployments fail due to an error when running the SQL upgrade script. After further investigation, you realize that this is because the SQL script is being executed on all of your Amazon EC2 Instances.

How would you ensure that the SQL script Is only executed once per deployment regardless of how many Amazon EC2 instances are running at the time?

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.

Use a container command within an Elastic Beanstalk configuration file to execute the script., ensuring that the leader on flag Is set to true.

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.

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.