


DOP-C01 Dumps

AWS Certified DevOps Engineer- Professional

<https://www.certleader.com/DOP-C01-dumps.html>



NEW QUESTION 1

You have decided that you need to change the instance type of your production instances which are running as part of an AutoScaling group. The entire architecture is deployed using CloudFormation Template. You currently have 4 instances in Production. You cannot have any interruption in service and need to ensure 2 instances are always running during the update? Which of the options below listed can be used for this?

- A. AutoScalingRollingUpdate
- B. AutoScalingScheduledAction
- C. AutoScalingReplacingUpdate
- D. AutoScalingIntegrationUpdate

Answer: A

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 Cloud Formation 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 Autoscaling updates, please refer to the below link: <https://aws.amazon.com/premiumsupport/knowledge-center/auto-scaling-group-rolling-updates/>

NEW QUESTION 2

You have an application running a specific process that is critical to the application's functionality, and have added the health check process to your Auto Scaling Group. The instances are showing healthy but the application itself is not working as it should. What could be the issue with the health check, since it is still showing the instances as healthy.

- 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 the application
- C. The health check is not configured properly
- D. The health check is not checking the application process

Answer: D

Explanation:

If you have custom health checks, you can send the information from your health checks to Auto Scaling so that Auto Scaling can use 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. For more information on Autoscaling health checks, please refer to the below document link: from AWS <http://docs.aws.amazon.com/autoscaling/latest/userguide/healthcheck.html>

NEW QUESTION 3

You have deployed an application to AWS which makes use of Autoscaling to launch new instances. You now want to change the instance type for the new instances. Which of the following is one of the action items to achieve this deployment?

- A. Use Elastic Beanstalk to deploy the new application with the new instance type
- B. Use Cloudformation to deploy the new application with the new instance type
- C. Create a new launch configuration with the new instance type
- D. Create new EC2 instances with the new instance type and attach it to the Autoscaling Group

Answer: C

Explanation:

The ideal way is to create a new launch configuration, attach it to the existing Auto Scaling group, and terminate the running instances. Option A is invalid because Elastic beanstalk cannot launch new instances on demand. Since the current scenario requires Autoscaling, this is not the ideal option. Option B is invalid because this will be a maintenance overhead, since you just have an Autoscaling Group. There is no need to create a whole Cloudformation template for this. Option D is invalid because Autoscaling Group will still launch EC2 instances with the older launch configuration. For more information on Autoscaling Launch configuration, please refer to the below document link: from AWS <http://docs.aws.amazon.com/autoscaling/latest/userguide/LaunchConfiguration.html>

NEW QUESTION 4

Which Auto Scaling process would be helpful when testing new instances before sending traffic to them, while still keeping them in your Auto Scaling Group?

- A. Suspend the process AZ Rebalance
- B. Suspend the process Health Check
- C. Suspend the process Replace Unhealthy
- D. Suspend the process AddToLoadBalancer

Answer: D

Explanation:

If you suspend AddToLoadBalancer, Auto Scaling launches the instances but does not add them to the load balancer or target group. If you resume the AddToLoadBalancer process, Auto Scaling resumes adding instances to the load balancer or target group when they are launched. However, Auto Scaling does not add the instances that were launched while this process was suspended. You must register those instances manually. Option A is invalid because this just balances the number of EC2 instances in the group across the Availability Zones in the region. Option B is invalid because this just checks the health of the instances. Auto Scaling marks an instance as unhealthy if Amazon EC2 or Elastic Load Balancing

tells

Auto Scaling that the instance is unhealthy.

Option C is invalid because this process just terminates instances that are marked as unhealthy and later creates new instances to replace them.

For more information on process suspension, please refer to the below document link: from AWS <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-suspend-resume-processes.html>

NEW QUESTION 5

Your company has multiple applications running on AWS. Your company wants to develop a tool that notifies on-call teams immediately via email when an alarm is triggered in your environment. You have multiple on-call teams that work different shifts, and the tool should handle notifying the correct teams at the correct times. How should you implement this solution?

- A. Create an Amazon SNS topic and an Amazon SQS queue
- B. 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
- C. Create an Amazon EC2 Auto Scaling group with both minimum and desired Instances configured to 0. Worker nodes in this group spawn when messages are added to the queue
- D. Workers then use Amazon Simple Email Service to send messages to your on-call teams.
- E. Create an Amazon SNS topic and configure your on-call team email addresses as subscriber
- F. Use the AWS SDK tools to integrate your application with Amazon SNS and send messages to this new topic
- G. Notifications will be sent to on-call users when a CloudWatch alarm is triggered.
- H. Create an Amazon SNS topic and configure your on-call team email addresses as subscriber
- I. Create a secondary Amazon SNS topic for alarms and configure your CloudWatch alarms to notify this topic when triggered
- J. Create an HTTP subscriber to this topic that notifies your application via HTTP POST when an alarm is triggered
- K. Use the AWS SDK tools to integrate your application with Amazon SNS and send messages to the first topic so that on-call engineers receive alerts.
- L. Create an Amazon SNS topic for each on-call group, and configure each of these with the team member emails as subscriber
- M. Create another Amazon SNS topic and configure your CloudWatch alarms to notify this topic when triggered
- N. Create an HTTP subscriber to this topic that notifies your application via HTTP POST when an alarm is triggered
- O. 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:

Option D fulfills all the requirements

1) First is to create a SNS topic for each group so that the required members get the email addresses.

2) Ensure the application uses the HTTPS endpoint and the SDK to publish messages. Option A is invalid because the SQS service is not required.

Option B and C are incorrect. As per the requirement we need to provide notification to only those on-call teams who are working in that particular shift when an alarm is triggered. It need not have to be sent to all the on-call teams of the company. With Option B & C, since we are not configuring the SNS topic for each on-call team the notifications will be sent to all the on-call teams. Hence these 2 options are invalid. For more information on setting up notifications, please refer to the below document link: from AWS http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/US_SetupSNS.html

NEW QUESTION 6

During metric analysis, your team has determined that the company's website during peak hours is experiencing response times higher than anticipated. You currently rely on Auto Scaling to make sure that you are scaling your environment during peak windows. How can you improve your Auto Scaling policy to reduce this high response time? Choose 2 answers.

- A. Push custom metrics to CloudWatch to monitor your CPU and network bandwidth from your servers, which will allow your Auto Scaling policy to 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: BD

Explanation:

Option B makes sense because maybe the max servers is low hence the application cannot handle the peak load.

Option D helps in ensuring Autoscaling can scale the group on the right metrics.

For more information on Autoscaling health checks, please refer to the below document link: from AWS

<http://docs.aws.amazon.com/autoscaling/latest/userguide/healthcheck.html>

NEW QUESTION 7

You have a code repository that uses Amazon S3 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.

- 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. Add an Amazon S3 bucket policy with a condition statement that requires multi-factor authentication in order to delete objects and enable bucket versioning.
- C. Use a configuration management service to deploy AWS Identity and Access Management user credentials to the Amazon EC2 instance
- D. Use these credentials to securely access the Amazon S3 bucket when deploying code.
- E. Create an Amazon Identity and Access Management role with authorization to access the Amazon S3 bucket, and launch all of your application's Amazon EC2 instances with this role.
- F. Use AWS Data Pipeline to lifecycle the data in your Amazon S3 bucket to Amazon Glacier on a weekly basis.
- G. Use AWS Data Pipeline with multi-factor authentication to securely deploy code from the Amazon S3 bucket to your Amazon EC2 instances.

Answer: BD

Explanation:

You can add another layer of protection by enabling MFA Delete on a versioned bucket. Once you do

so, you must provide your AWS account's access keys and a valid code from the account's MFA device in order to permanently delete an object version or suspend or reactivate versioning on the bucket. For more information on MFA please refer to the below link: <https://aws.amazon.com/blogs/security/securing-access-to-aws-using-mfa-part-3/> 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 For more information on Roles for EC2 please refer to the below link: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html> Option A is invalid because this will not address either the integrity or security concern completely. Option C is invalid because user credentials should never be used in EC2 instances to access AWS resources. Option D and F are invalid because AWS IAM is an unnecessary overhead when you already have inbuilt controls to manage security for S3.

NEW QUESTION 8

You are a DevOps 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

- A. Re-deploy your application using a CloudFormation template to deploy Elastic Beanstalk
- B. Re-deploy with a CloudFormation template, define update policies on Auto Scaling groups in your CloudFormation template
- C. Use UpdatePolicy attribute to specify how CloudFormation handles updates to Auto Scaling Group resource.
- D. After each stack is deployed, tear down the old stack

Answer: BC

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. Option A is invalid because it is not efficient to use CloudFormation to use Elastic Beanstalk. Option D is invalid because this is an inefficient process to tear down stacks when there are stack policies available. For more information on Autoscaling Rolling Updates please refer to the below link:
• <https://aws.amazon.com/premiumsupport/knowledge-center/auto-scaling-group-rolling-updates/>

NEW QUESTION 9

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

- A. Create an Amazon S3 bucket per user, and use your application to generate the S3 URI for the appropriate content.
- B. Use AWS Identity and Access Management (IAM) user accounts as your application-level user database, and offload the burden of authentication from your application code.
- C. Authenticate your users at the application level, and use AWS Security Token Service (STS) to grant token-based authorization to S3 objects.
- D. Authenticate your users at the application level, and send an SMS token message to the user
- E. Create an Amazon S3 bucket with the same name as the SMS message token, and move the user's objects to that bucket.
- F. Use a key-based naming scheme comprised from the user IDs for all user objects in a single Amazon S3 bucket.

Answer: CE

Explanation:

The AWS Security Token Service (STS) is a web service that enables you to request temporary, limited-privilege credentials for AWS Identity and Access Management (IAM) users or for users that you authenticate (federated users). The token can then be used to grant access to the objects in S3. You can then provide access to the objects based on the key values generated via the user id. Option A is possible but then becomes a maintenance overhead because of the number of buckets. Option B is invalid because IAM users is not a good security practice. Option D is invalid because SMS tokens are not efficient for this requirement. For more information on the Security Token Service please refer to the below link: <http://docs.aws.amazon.com/STS/latest/APIReference/Welcome.html>

NEW QUESTION 10

You have an Auto Scaling group with 2 AZs. One AZ has 4 EC2 instances and the other has 3 EC2 instances. None of the instances are protected from scale in. Based on the default Auto Scaling termination policy what will happen?

- A. Auto Scaling selects an instance to terminate randomly
- B. Auto Scaling will terminate unprotected instances in the Availability Zone with the oldest launch configuration.
- C. Auto Scaling terminates which unprotected instances are closest to the next billing hour.
- D. Auto Scaling will select the AZ with 4 EC2 instances and terminate an instance.

Answer: D

Explanation:

The default termination policy is designed to help ensure that your network architecture spans Availability Zones evenly. When using the default termination policy, Auto Scaling selects an instance to terminate as follows:
Auto Scaling determines whether there are instances in multiple Availability Zones. If so, it selects the Availability Zone with the most instances and at least one instance that is not protected from scale in. If there is more than one Availability Zone with this number of instances, Auto Scaling selects the Availability Zone with the instances that use the oldest launch configuration. For more information on Autoscaling instance termination please refer to the below link:
<http://docs.aws.amazon.com/autoscaling/latest/userguide/as-instance-termination.html>

NEW QUESTION 10

You work for a company that has multiple applications which are very different and built on different programming languages. How can you deploy applications as quickly as possible?

- A. Develop each app in one Docker container and deploy using ElasticBeanstalk

- B. Create a Lambda function deployment package consisting of code and any dependencies
- C. Develop each app in a separate Docker container and deploy using Elastic Beanstalk V
- D. Develop each app in a separate Docker containers and deploy using CloudFormation

Answer: C

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 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.

Option A is an efficient way to use Docker. The entire idea of Docker is that you have a separate environment for various applications.

Option B is ideally used to running code and not packaging the applications and dependencies Option D is not ideal deploying Docker containers using Cloudformation

For more information on Docker and Clastic Beanstalk, please visit the below URL:

? [http://docs.aws.a mazon.com/elasticbeanstalk/latest/dg/create_deploy_docker.html](http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker.html)

NEW QUESTION 13

You have a current Clouformation template defines in AWS. You need to change the current alarm threshold defined in the Cloudwatch alarm. How can you achieve this?

- A. Currently, there is no option to change what is already defined in Cloudformation templates.
- B. Update the template and then update the stack with the new templat
- C. Automatically all resources will be changed in the stack.
- D. Update the template and then update the stack with the new templat
- E. Only those resources that need to be changed will be change
- F. All other resources which do not need to be changed will remain as they are.
- G. Delete the current cloudformation templat
- H. Create a new one which will update the current resources.

Answer: C

Explanation:

Option A is incorrect because Cloudformation templates have the option to update resources.

Option B is incorrect because only those resources that need to be changed as part of the stack update are actually updated.

Option D is incorrect because deleting the stack is not the ideal option when you already have a change option available.

When you need to make changes to a stack's settings or change its resources, you update the stack instead of deleting it and creating a new stack. For example, if you

have a stack with an EC2 instance, you can update the stack to change the instance's AMI ID.

When you update a stack, you submit changes, such as new input parameter values or an updated template. AWS CloudFormation compares the changes you submit with the current state of your stack and updates only the changed resources

For more information on stack updates please refer to the below link:

• [http://docs.aws.a mazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-updating- stacks.html](http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-updating-stacks.html)

NEW QUESTION 15

After reviewing the last quarter's monthly bills, management has noticed an increase in the overall bill from Amazon. After researching this increase in cost, you discovered that one of your new services is doing a lot of GET Bucket API calls to Amazon S3 to build a metadata cache of all objects in the applications bucket. Your boss has asked you to come up with a new cost-effective way to help reduce the amount of these new GET Bucket API calls. What process should you use to help mitigate the cost?

- A. Update your Amazon S3 buckets' lifecycle policies to automatically push a list of objects to a new bucket, and use this list to view objects associated with the application's bucket.
- B. Create a new DynamoDB tabl
- C. Use the new DynamoDB table to store all metadata about all objects uploaded to Amazon S3. Any time a new object is uploaded, update the application's internalAmazon S3 object metadata cache from DynamoDB.C Using Amazon SNS, create a notification on any new Amazon S3 objects that automatical ly updates a new DynamoDB table to store allmetadata about the new objec
- D. Subscribe the application to the Amazon SNS topic to update its internal Amazon S3 object metadata cache from the DynamoDB tabl
- E. ^/
- F. Upload all files to an ElastiCache file cache serve
- G. Update your application to now read all file metadata from the ElastiCache file cache server, and configure the ElastiCache policies to push all files to Amazon S3 for long-term storage.

Answer: C

Explanation:

Option A is an invalid option since Lifecycle policies are normally used for expiration of objects or archival of objects.

Option B is partially correct where you store the data in DynamoDB, but then the number of GET requests would still be high if the entire DynamoDB table had to be

traversed and each object compared and updated in S3.

Option D is invalid because uploading all files to Clastic Cache is not an ideal solution.

The best option is to have a notification which can then trigger an update to the application to update the DynamoDB table accordingly.

For more information on SNS triggers and DynamoDB please refer to the below link:

? <https://aws.amazon.com/blogs/compute/619/>

NEW QUESTION 18

You have a complex system that involves networking, IAM policies, and multiple, three-tier 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 CloudFormation to define these AWS resources so that you can automate and version-control your infrastructure. How would you use AWS CloudFormation to provide agile new environments for your customers in a cost-effective, reliable manner?

- 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, create nested stacks in AWS CloudFormation, and maintain several templates to version-control.
- C. ➤/
- D. 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.
- E. Manually construct the networking layer using Amazon Virtual Private Cloud (VPC) because this does not change often, and then use AWS CloudFormation to define all other ephemeral resources.

Answer: B

Explanation:

As your infrastructure grows, common patterns can emerge in which you declare the same components in each of your templates. You can separate out these common components and create dedicated templates for them. That way, you can mix and match different templates but use nested stacks to create a single, unified stack. Nested stacks are stacks that create other stacks. To create nested stacks, use the `AWS::CloudFormation::StackResource` in your template to reference other templates.

For more information on CloudFormation best practises please refer to the below link: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html>

NEW QUESTION 20

You are using CloudFormation to launch an EC2 instance and then configure an application after the instance is launched. You need the stack creation of the ELB and Auto Scaling to wait until the EC2 instance is launched and configured properly. How do you do this?

- A. It is not possible for the stack creation to wait until one service is created and launched
- B. Use the `WaitCondition` resource to hold the creation of the other dependent resources
- C. Use a `CreationPolicy` to wait for the creation of the other dependent resources ➤/
- D. Use the `HoldCondition` resource to hold the creation of the other dependent resources

Answer: C

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'll know your applications are ready to go after stack creation succeeds.

A `CreationPolicy` instructs CloudFormation to wait on an instance until CloudFormation receives the specified number of signals

Option A is invalid because this is possible

Option B is invalid because this is used to make AWS CloudFormation pause the creation of a stack and wait for a signal before it continues to create the stack

For more information on this, please visit the below URL:

- <https://aws.amazon.com/blogs/devops/use-a-creationpolicy-to-wait-for-on-instance-configurations/>

NEW QUESTION 21

You have enabled Elastic Load Balancing HTTP health checking. After looking at the AWS Management Console, you see that all instances are passing health checks, but your customers are reporting that your site is not responding. What is the cause?

- A. The HTTP health checking system is misreporting due to latency in inter-instance metadata synchronization.
- B. The health check in place is not sufficiently evaluating the application function.
- C. The application is returning a positive health check too quickly for the AWS Management Console to respond.
- D- Latency in DNS resolution is interfering with Amazon EC2 metadata retrieval.

Answer: B

Explanation:

You need to have a custom health check which will evaluate the application functionality. It's not enough using the normal health checks. If the application functionality does not work and if you don't have custom health checks, the instances will still be deemed as healthy.

If you have custom health checks, you can send the information from your health checks to Auto Scaling so that Auto Scaling can use 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

For more information on Autoscaling health checks, please refer to the below document link: from AWS

<http://docs.aws.amazon.com/autoscaling/latest/userguide/healthcheck.html>

NEW QUESTION 25

Your development team wants account-level access to production instances in order to do live debugging of a highly secure environment. Which of the following should you do?

- A. Place the credentials provided by Amazon Elastic Compute Cloud (EC2) into a secure Amazon Simple Storage Service (S3) bucket with encryption enabled
- B. Assign AWS Identity and Access Management (IAM) users to each developer so they can download the credentials file.
- 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 each developer's own public key into a private S3 bucket, use instance profiles and configuration management to create a user account for each developer on all instances, and place the user's public keys into the appropriate account
- E. ^/
- F. 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.

Answer: C

Explanation:

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.

Option A and D are invalid, because the credentials should not be provided by a AWS EC2 Instance. Option B is invalid because you would not create a service account, instead you should create an instance profile.

For more information on Instance profiles, please refer to the below document link: from AWS

- http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_switch-role-ec2-instance-profiles.html

NEW QUESTION 28

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?

- A. Develop the applications in Docker containers, and then deploy them to Elastic Beanstalk environments with Auto Scaling and Elastic Load Balancing.
- B. Use the AWS CloudFormation Docker import service to build and deploy the applications with high availability in multiple Availability Zones.
- C. 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.
- D. 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 multiple Availability Zones.

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 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.

By using Docker with Elastic Beanstalk, you have an infrastructure that automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring.

For more information on Dockers and Elastic beanstalk please refer to the below link:

- http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker.html

NEW QUESTION 33

Your application uses Cloud Formation to orchestrate your application's resources. During your testing phase before the application went live, your Amazon RDS instance type was changed and caused the instance to be re-created, resulting in the loss of test data. How should you prevent this from occurring in the future?

- A. Within the AWS CloudFormation parameter with which users can select the Amazon RDS instance type, set AllowedValues to only contain the current instance type.
- B. Use an AWS CloudFormation stack policy to deny updates to the instance.
- C. Only allow UpdateStack permission to IAM principals that are denied SetStackPolicy.
- D. In the AWS CloudFormation template, set the AWS::RDS::DBInstance's DBInstanceClass property to be read-only.
- E. Subscribe to the AWS CloudFormation notification "BeforeResourceUpdate," and call CancelStackUpdate if the resource identified is the Amazon RDS instance.
- F. Update the stack using ChangeSets

Answer: E

Explanation:

When you need to update a stack, understanding how your changes will affect running resources before you implement them can help you update stacks with confidence. Change sets allow you to preview how proposed changes to a stack might impact your running resources, for example, whether your changes will delete or replace any critical resources. AWS CloudFormation makes the changes to your stack only when you decide to execute the change set, allowing you to decide whether to proceed with your proposed changes or explore other changes by creating another change set.

For example, you can use a change set to verify that AWS CloudFormation won't replace your stack's database instances during an update.

NEW QUESTION 35

As an architect you have decided to use CloudFormation instead of OpsWorks or Elastic Beanstalk for deploying the applications in your company. Unfortunately, you have discovered that there is a resource type that is not supported by CloudFormation. What can you do to get around this.

- A. Specify more mappings and separate your template into multiple templates by using nested stacks.
- B. Create a custom resource type using template developer, custom resource template, and CloudFormation.
- C. */
- D. Specify the custom resource by separating your template into multiple templates by using nested stacks.
- E. Use a configuration management tool such as Chef, Puppet, or Ansible.

Answer: B

Explanation:

Custom resources enable you to write custom provisioning logic in templates that AWS CloudFormation runs anytime you create, update (if you changed the custom resource), or delete stacks. For example, you might want to include resources that aren't available as AWS CloudFormation resource types. You can include those resources by using custom resources. That way you can still manage all your related resources in a single stack.

For more information on custom resources in CloudFormation please visit the below URL:

- <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/template-custom-resources.html>

NEW QUESTION 40

You work for a startup that has developed a new photo-sharing application for mobile devices. Over recent months your application has increased in popularity; this has resulted in a decrease in the performance of the application due to the increased load. Your application has a two-tier architecture that is composed of an Auto Scaling PHP application tier and a MySQL RDS instance initially deployed with AWS CloudFormation. Your Auto Scaling group has a min value of 4 and a max value of 8. The desired capacity is now at 8 because of the high CPU utilization of the instances. After some analysis, you are confident that the performance issues stem from a constraint in CPU capacity, although memory utilization remains low. You therefore decide to move from the general-purpose M3 instances to the compute-optimized C3 instances. How would you deploy this change while minimizing any interruption to your end users?

- A. Sign into the AWS Management Console, copy the old launch configuration, and create a new launch configuration that specifies the C3 instance.
- B. Update the Auto Scaling group with the new launch configuration.

- C. 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 typ
- E. Add an UpdatePolicy attribute to your Auto Scaling group that specifies AutoScalingRollingUpdate.
- F. Update the launch configuration specified in the AWS CloudFormation template with the new C3 instance typ
- G. Run a stack update with the new templat
- H. Auto Scaling will then update the instances with the new instance type.
- I. Update the launch configuration specified in the AWS CloudFormation template with the new C3instance typ
- J. Also add an UpdatePolicy attribute to your Auto Scalinggroup that specifies AutoScalingRollingUpdat
- K. Run a stack update with the new template.

Answer: D

Explanation:

The AWS::AutoScaling::AutoScalingGroup resource supports an UpdatePolicy attribute. This is used to define how an Auto Scalinggroup resource is updated when an update to the Cloud Formation 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 rolling updates, please visit the below link:

- <https://aws.amazon.com/premiumsupport/knowledge-center/auto-scaling-group-rolling-updates/>

NEW QUESTION 43

You've been tasked with improving the current deployment process by making it easier to deploy and reducing the time it takes. You have been tasked with creating a continuous integration (CI) pipeline that can build AMI'S. Which of the below is the best manner to get this done. Assume that at max your development team will be deploying builds 5 times a week.

- A. Use a dedicated EC2 instance with an EBS Volum
- B. Download and configure the code and then crate an AMI out of that.
- C. Use OpsWorks to launch an EBS-backed instance, then use a recipe to bootstrap the instance, and then have the CI system use the CreateImage API call to make an AMI from it.
- D. Upload the code and dependencies to Amazon S3, launch an instance, download the package fromAmazon S3, then create the AMI with the CreateSnapshot API call
- E. Have the CI system launch a new instance, then bootstrap the code and dependencies on that instance, and create an AMI using the CreateImage API call.

Answer: D

Explanation:

Since the number of calls is just a few times a week, there are many open source systems such as Jenkins which can be used as CI based systems.

Jenkins can be used as an extensible automation server, Jenkins can be used as a simple CI server or turned into the continuous delivery hub for any project.

For more information on the Jenkins CI tool please refer to the below link:

- <https://jenkins.io/>

Option A and C are partially correct, but since you just have 5 deployments per week, having separate instances which consume costs is not required. Option B is partially correct, but again having a separate system such as Opswork for such a low number of deployments is not required.

NEW QUESTION 46

You have deployed an Elastic Beanstalk application in a new environment and want to save the current state of your environment in a document. You want to be able to restore your environment to the current state later or possibly create a new environment. You also want to make sure you have a restore point. How can you achieve this?

- A. Use CloudFormation templates
- B. Configuration Management Templates
- C. Saved Configurations
- D. Saved Templates

Answer: C

Explanation:

You can save your environment's configuration as an object in Amazon S3 that can be applied to other environments during environment creation, or applied to a running environment. Saved configurations are YAML formatted templates that define an environment's platform configuration, tier, configuration option settings, and tags.

For more information on Saved Configurations please refer to the below link:

- <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/environment-configuration-savedconfig.html>

NEW QUESTION 49

You need to monitor specific metrics from your application and send real-time alerts to your Devops Engineer. Which of the below services will fulfil this requirement? Choose two answers

- A. Amazon CloudWatch
- B. Amazon Simple Notification Service
- C. Amazon Simple Queue Service
- D. Amazon Simple Email Service

Answer: AB

Explanation:

Amazon Cloud Watch monitors your Amazon Web Services (AWS) resources and the applications you run on AWS in real time. You can use Cloud Watch to collect and track metrics, which are variables you can measure for your resources and applications. Cloud Watch alarms send notifications or automatically make changes to the resources you are monitoring based on rules that you define.

For more information on AWS Cloudwatch, please refer to the below document link: from AWS

- <http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/WhatIsCloudWatch.htm> | Amazon Cloud Watch uses Amazon SNS to send email. First, create and subscribe to an SNS topic.

When you create a CloudWatch alarm, you can add this SNS topic to send an email notification when the alarm changes state

For more information on AWS Cloudwatch and SNS, please refer to the below document link: from AWS
http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/US_SetupSNS.html

NEW QUESTION 50

Your company has developed a web application and is hosting it in an Amazon S3 bucket configured for static website hosting. The application is using the AWS SDK for JavaScript in the browser to access data stored in an Amazon DynamoDB table. How can you ensure that API keys for access to your data in DynamoDB are kept secure?

- A. Create an Amazon S3 role in 1AM with access to the specific DynamoDB tables, and assign it to the bucket hosting your website.
- B. Configure S3 bucket tags with your AWS access keys for your bucket hosing your website so that the application can query them for access.
- C. Configure a web identity federation role within 1AM to enable access to the correct DynamoDB resources and retrieve temporary credentials.
- D. Store AWS keys in global variables within your application and configure the application to use these credentials when making requests.

Answer: C

Explanation:

With web identity federation, you don't need to create custom sign-in code or manage your own user identities. Instead, users of your app can sign in using a well-known identity provider (IdP) — such as Login with Amazon, Facebook, Google, or any other OpenID Connect (OIDC)-compatible IdP, receive an authentication token, and then exchange that token for temporary security credentials in AWS that map to an 1AM role with permissions to use the resources in your AWS account. Using an IdP helps you keep your AWS account secure, because you don't have to embed and distribute long- term security credentials with your application. For more information on Web Identity Federation, please refer to the below document link: from AWS
http://docs.wsamazon.com/IAM/latest/UserGuide/id_roles_providers_oidc.html

NEW QUESTION 51

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 MYSQL DB instance. How should you securely store these credentials?

- A. Give the Amazon EC2 instances an 1AM role that allows read access to a private Amazon S3 bucke
- B. Store a file with database credentials in the Amazon S3 bucke
- C. Have your configuration management system pull the file from the bucket when it is needed.
- D. Launch an Amazon EC2 instance and use the configuration management system to bootstrap the instance with the Amazon RDS DB credential
- E. Create an AML from this instance.
- F. Store the Amazon RDS DB credentials in Amazon EC2 user dat
- G. Import the credentials into the Instance on boot.
- H. Assign an 1AM role to your Amazon EC2 instance, and use this 1AM role to access the Amazon RDS DB from your Amazon EC2 instances.

Answer: D

Explanation:

Creating and Using an 1AM Policy for 1AM Database Access

To allow an 1AM user or role to connect to your DB instance or DB cluster, you must create an 1AM policy. After that you attach the policy to an 1AM user or role.
Note

To learn more about 1AM policies, see Authentication and Access Control for Amazon RDS.

The following example policy allows an 1AM user to connect to a DB instance using 1AM database authentication.



Important

Don't confuse the rds-db: prefix with other Amazon RDS action prefixes that begin with rds:. You use the rds-db: prefix and the rds-db:connect action only for 1AM database authentication. They aren't valid in any other context.

1AM Database Authentication for MySQL and Amazon Aurora

With Amazon RDS for MySQL or Aurora with MySQL compatibility, you can authenticate to your DB instance or DB cluster using AWS Identity and Access Management (IAM) database authentication. With this authentication method, you don't need to use a password when you connect to a DB instance. Instead, you use an authentication token.

An authentication token is a unique string of characters that Amazon RDS generates on request. Authentication tokens are generated using AWS Signature Version 4. Each token has a lifetime of 15 minutes. You don't need to store user credentials in the database, because authentication is managed externally using 1AM. You can also still use standard database authentication.

IAM database authentication provides the following benefits:

- Network traffic to and from the database is encrypted using Secure Sockets Layer (SSL).
- You can use IAM to centrally manage access to your database resources, instead of managing access individually on each DB instance or DB cluster.
- For applications running on Amazon EC2, you can use EC2 instance profile credentials to access the database instead of a password, for greater security.

For more information please refer to the below document link from AWS

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/UsingWithRDS.IAMDBAuth.html>

<https://docs^ws.amazon.com/AmazonRDS/latest/UserGuide/UsingWithRDS.IAMDBAuth.IAMPolicy.html>

You can use roles to delegate access to users, applications, or services that don't normally have access to your AWS resources. For example, you might want to grant users in your AWS account access to resources they don't usually have, or grant users in one AWS account access to resources in another account. Or you might want to allow a mobile app to use AWS resources, but not want to embed AWS keys within the app (where they can be difficult to rotate and where users can potentially extract them). Sometimes you want to give AWS access to users who already have identities defined outside of AWS, such as in your corporate directory. Or, you might want to grant access to your account to third parties so that they can perform an audit on your resources. For more information on 1AM Roles, please refer to the below document link: from AWS

http://docs.aws.a mazon.com/IAM/latest/UserGuide/id_roles.html

NEW QUESTION 54

What is web identity federation?

- A. Use of an identity provider like Google or Facebook to become an AWS IAM User.
- B. Use of an identity provider like Google or Facebook to exchange for temporary AWS security credentials.
- C. Use of AWS IAM User tokens to log in as a Google or Facebook user.
- D. Use STS service to create an user on AWS which will allow them to login from facebook or google app.

Answer: B

Explanation:

With web identity federation, you don't need to create custom sign-in code or manage your own user identities. Instead, users of your app can sign in using a well-known identity provider (IdP) — such as Login with Amazon, Facebook, Google, or any other OpenID Connect (OIDC)-compatible IdP, receive an authentication token, and 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 don't have to embed and distribute long-term security credentials with your application. For more information on Web Identity federation please refer to the below link:
http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_oidc.html

NEW QUESTION 58

You have decided to migrate your application to the cloud. You cannot afford any downtime. You want to gradually migrate so that you can test the application with a small percentage of users and increase over time. Which of these options should you implement?

- A. Use Direct Connect to route traffic to the on-premise location
- B. In DirectConnect, configure the amount of traffic to be routed to the on-premise location.
- C. Implement a Route 53 failover routing policy that sends traffic back to the on-premises application if the AWS application fails.
- D. Configure an Elastic Load Balancer to distribute the traffic between the on-premises application and the AWS application.
- E. Implement a Route 53 weighted routing policy that distributes the traffic between your on-premises application and the AWS application depending on weight.

Answer: D

Explanation:

Option A is incorrect because DirectConnect cannot control the flow of traffic.

Option B is incorrect because you want to split the percentage of traffic. Failover will direct all of the traffic to the backup servers.

Option C is incorrect because you cannot control the percentage distribution of traffic.

Weighted routing lets you associate multiple resources with a single domain name (example.com) or subdomain name (acme.example.com) and choose how much traffic is routed to each resource. This can be useful for a variety of purposes, including load balancing and testing new versions of software.

For more information on the Routing policy please refer to the below link: <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html>

NEW QUESTION 63

You are hired as the new head of operations for a SaaS company. Your CTO has asked you to make debugging any part of your entire operation simpler and as fast as possible. She complains that she has no idea what is going on in the complex, service-oriented architecture, because the developers just log to disk, and it's very hard to find errors in logs on so many services. How can you best meet this requirement and satisfy your CTO?

- A. Copy all log files into AWS S3 using a cron job on each instance
- B. Use an S3 Notification Configuration on the PutBucket event and publish events to AWS Lambda
- C. Use the Lambda to analyze logs as soon as they come in and flag issues.
- D. Begin using CloudWatch Logs on every service
- E. Stream all Log Groups into S3 object
- F. Use AWS EMR cluster jobs to perform ad-hoc MapReduce analysis and write new queries when needed.
- G. Copy all log files into AWS S3 using a cron job on each instance
- H. Use an S3 Notification Configuration on the PutBucket event and publish events to AWS Kinesis
- I. Use Apache Spark on AWS EMR to perform at-scale stream processing queries on the log chunks and flag issues.
- J. Begin using CloudWatch Logs on every service
- K. Stream all Log Groups into an AWS Elasticsearch Service Domain running Kibana 4 and perform log analysis on a search cluster.

Answer: D

Explanation:

Amazon Elasticsearch Service makes it easy to deploy, operate, and scale Elasticsearch for log analytics, full text search, application monitoring, and more.

Elasticsearch Service is a fully managed service that delivers Elasticsearch's easy-to-use APIs and real-time capabilities along with the availability, scalability, and security required by production workloads. The service offers built-in integrations with Kibana, Logstash, and AWS services including Amazon Kinesis Firehose, AWS Lambda, and Amazon Cloud Watch so that you can go from raw data to actionable insights quickly. For more information on Elastic Search, please refer to the below link:

- <https://aws.amazon.com/elasticsearch-service/>

NEW QUESTION 66

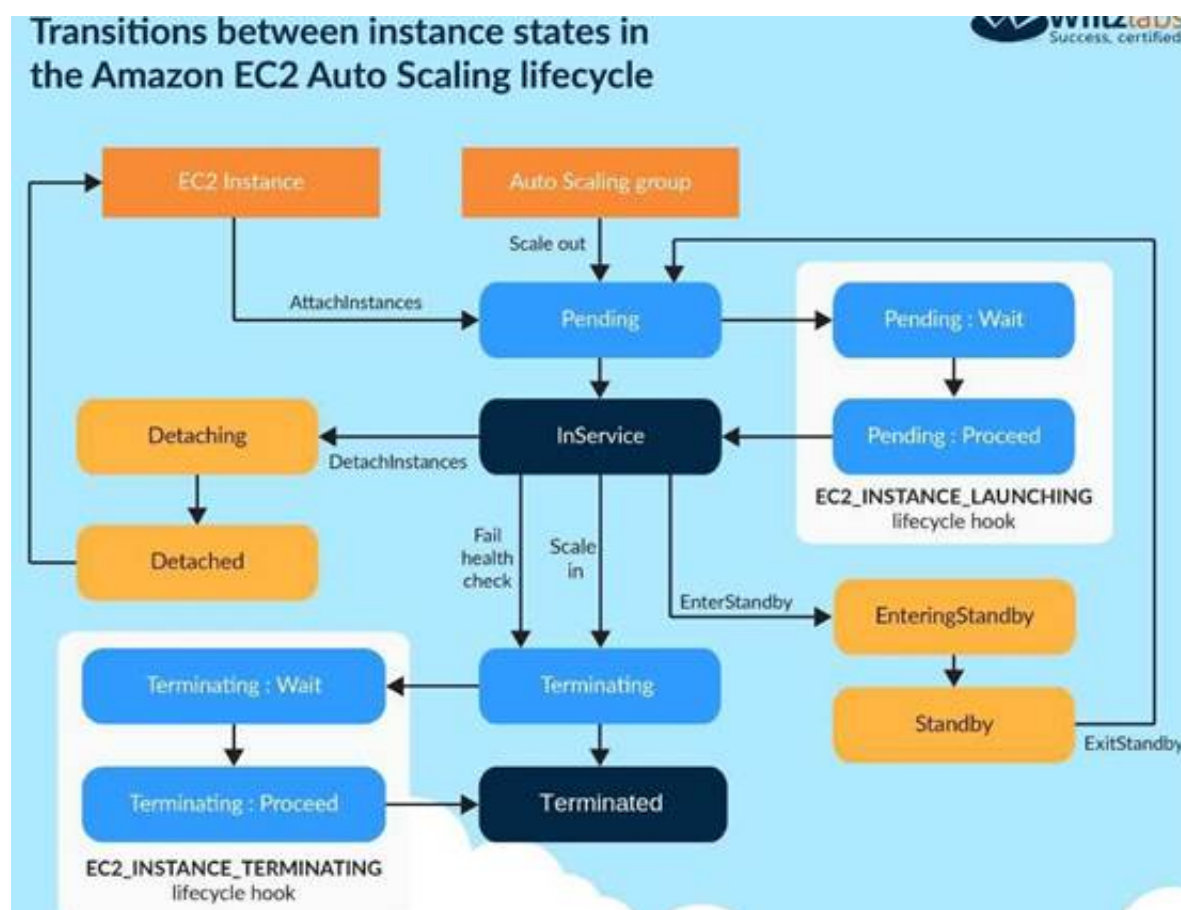
For AWS Auto Scaling, what is the first transition state an instance enters after leaving steady state when scaling in due to health check failure or decreased load?

- A. Terminating
- B. Detaching
- C. Terminating:Wait
- D. EnteringStandby

Answer: A

Explanation:

The below diagram shows the Lifecycle policy. When the scale-in happens, the first action is the Terminating action.



For more information on Autoscaling Lifecycle, please refer to the below link:
<http://docs.aws.amazon.com/autoscaling/latest/userguide/AutoScalingGroupLifecycle.html>

NEW QUESTION 67

You have an application hosted in AWS. You wanted to ensure that when certain thresholds are reached, a DevOps Engineer is notified. Choose 3 answers from the options given below

- A. Use CloudWatch Logs agent to send log data from the app to CloudWatch Logs from Amazon EC2 instances
- B. Pipe data from EC2 to the application logs using AWS Data Pipeline and CloudWatch
- C. Once a CloudWatch alarm is triggered, use SNS to notify the Senior DevOps Engineer.
- D. Set the threshold your application can tolerate in a CloudWatch Logs group and link a CloudWatch alarm on that threshold.

Answer: ACD

Explanation:

You can use Cloud Watch Logs to monitor applications and systems using log data. For example, CloudWatch Logs can track the number of errors that occur in your application logs and send you a notification whenever the rate of errors exceeds a threshold you specify. CloudWatch Logs uses your log data for monitoring; so, no code changes are required. For example, you can monitor application logs for specific literal terms (such as "NullPointerException") or count the number of occurrences of a literal term at a particular position in log data (such as "404" status codes in an Apache access log). When the term you are searching for is found, CloudWatch Logs reports the data to a CloudWatch metric that you specify. For more information on Cloudwatch Logs please refer to the below link:
<http://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html>
Amazon CloudWatch uses Amazon SNS to send email. First, create and subscribe to an SNS topic. When you create a CloudWatch alarm, you can add this SNS topic to send an email notification when the alarm changes state. For more information on Cloudwatch and SNS please refer to the below link:
http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/US_SetupSNS.html

NEW QUESTION 68

Your company releases new features with high frequency while demanding high application availability. As part of the application's A/B testing, logs from each updated Amazon EC2 instance of the application need to be analyzed in near real-time, to ensure that the application is working flawlessly after each deployment. If the logs show any anomalous behavior, then the application version of the instance is changed to a more stable one. Which of the following methods should you use for shipping and analyzing the logs in a highly available manner?

- A. Ship the logs to Amazon S3 for durability and use Amazon EMR to analyze the logs in a batch manner each hour.
- 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 an Amazon Kinesis stream and have the consumers analyze the logs in a live manner.
- D. Ship the logs to a large Amazon EC2 instance and analyze the logs in a live manner.

Answer: C

Explanation:

Answer - C
You can use Kinesis Streams for rapid and continuous data intake and aggregation. The type of data used includes IT infrastructure log data, application logs, social media, market data feeds, and web clickstream data. Because the response time for the data intake and processing is in real time, the processing is typically lightweight. The following are typical scenarios for using Kinesis Streams:
• Accelerated log and data feed intake and processing - You can have producers push data directly into a stream. For example, push system and application logs and they'll be available for processing in seconds. This prevents the log data from being lost if the front end or application server fails. Kinesis Streams provides accelerated data feed intake because you don't batch the data on the servers before you submit it for intake.
• 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 wait to receive batches of data. For more information on Amazon Kinesis and SNS please refer to the below link:
• <http://docs.aws.amazon.com/kinesis/latest/dev/introduction.html>

NEW QUESTION 73

You have an application running on an Amazon EC2 instance and you are using IAM roles to securely access AWS Service APIs. How can you configure your application running on that instance to retrieve the API keys for use with the AWS SDKs?

- A. 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.
- B. Within your application code, make a GET request to the IAM Service API to retrieve credentials for your user.
- C. 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.
- D. 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.

Answer: 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

For more information on Roles for EC2 please refer to the below link: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

NEW QUESTION 75

You need to deploy a Node.js application and do not have any experience with AWS. Which deployment method will be the simplest for you to deploy?

- A. AWS Elastic Beanstalk
- B. AWS CloudFormation
- C. AWS EC2
- D. AWS OpsWorks

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

For more information on Elastic beanstalk please refer to the below link:

- <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/Welcome.html>

NEW QUESTION 79

You have been given a business requirement to retain log files for your application for 10 years. You need to regularly retrieve the most recent logs for troubleshooting. Your logging system must be cost-effective, given the large volume of logs. What technique should you use to meet these requirements?

- A. Store your log in Amazon CloudWatch Logs.
- B. Store your logs in Amazon Glacier.
- C. Store your logs in Amazon S3, and use lifecycle policies to archive to Amazon Glacier.
- D. Store your logs on Amazon EBS, and use Amazon EBS snapshots to archive them.

Answer: C

Explanation:

Option A is invalid, because CloudWatch will not store the logs indefinitely and secondly it won't be the cost-effective option.

Option B is invalid, because it won't serve the purpose of regularly retrieving the most recent logs for troubleshooting. You will need to pay more to retrieve the logs faster from this storage.

Option D is invalid, because it is not an ideal or cost-effective option.

You can define lifecycle configuration rules for objects that have a well-defined lifecycle. For example: if you are uploading periodic logs to your bucket, your application might need these logs for a week or a month after creation, and after that you might want to delete them.

Some documents are frequently accessed for a limited period of time. After that, these documents are less frequently accessed. Over time, you might not need real-time access to these objects, but your organization or regulations might require you to archive them for a longer period and then optionally delete them later.

You might also upload some types of data to Amazon S3 primarily for archival purposes, for example digital media archives, financial and healthcare records, raw genomics sequence data, long-term database backups, and data that must be retained for regulatory compliance.

For more information on Lifecycle management please refer to the below link: <http://docs.aws.amazon.com/AmazonS3/latest/dev/object-lifecycle-mgmt.html>

Note:

Option C is the cheapest option, but CloudWatch can store logs indefinitely or between 10 years and one day.

"Log Retention—By default, logs are kept indefinitely and never expire. You can adjust the retention policy for each log group, keeping the indefinite retention, or choosing a retention period between 10 years and one day." <https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html>

NEW QUESTION 82

You have deployed a CloudFormation template which is used to spin up resources in your account. Which of the following status in CloudFormation represents a failure?

- A. UPDATE_COMPLETE_CLEANUP_IN_PROGRESS
- B. DELETE_COMPLETE
- C. ROLLBACK_IN_PROGRESS
- D. UPDATE_IN_PROGRESS

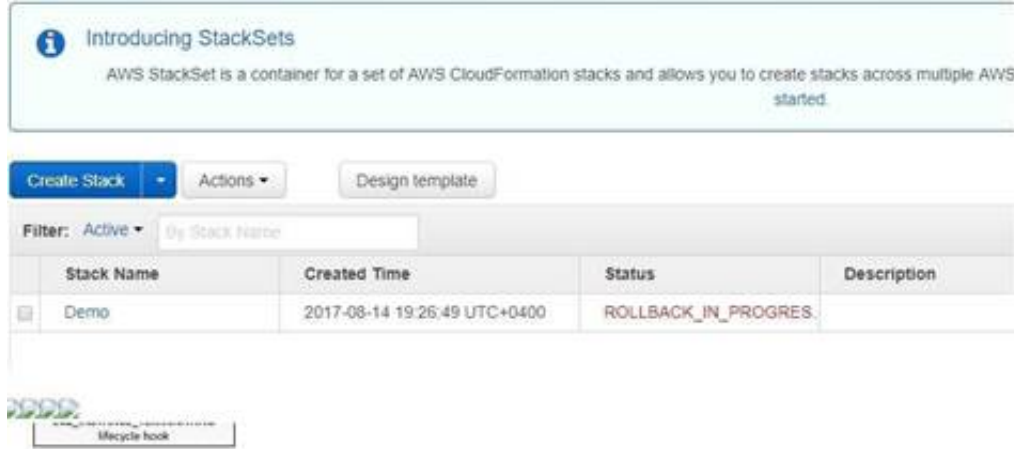
Answer: C

Explanation:

AWS CloudFormation provisions and configures resources by making calls to the AWS services that are described in your template.

After all the resources have been created, AWS CloudFormation reports that your stack has been created. You can then start using the resources in your stack. If

stack creation fails, AWS CloudFormation rolls back your changes by deleting the resources that it created. The below snapshot from Cloudformation shows what happens when there is an error in the stack creation.



For more information on how Cloud Formation works, please refer to the below link: <http://docs.ws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-what-is-howdoesitwork.html>

NEW QUESTION 85

You need to perform ad-hoc analysis on log data, including searching quickly for specific error codes and reference numbers. Which should you evaluate first?

- A. AWS Elasticsearch Service
- B. AWSRedShift
- C. AWSEMR
- D. AWS DynamoDB

Answer: A

Explanation:

Amazon Dasticsearch Service makes it easy to deploy, operate, and scale dasticsearch for log analytics, full text search, application monitoring, and more. Amazon

Oasticsearch Service is a fully managed service that delivers Dasticsearch's easy-to-use APIs and real- time capabilities along with the availability, scalability, and security required by production workloads. The service offers built-in integrations with Kibana, Logstash, and AWS services including Amazon Kinesis Firehose, AWS Lambda, and Amazon CloudWatch so that you can go from raw data to actionable insights quickly For more information on the elastic cache service, please refer to the below link:

- <https://aws.amazon.com/elasticsearch-service/>

NEW QUESTION 88

You have a requirement to host a cluster of NoSQL databases. There is an expectation that there will be a lot of I/O on these databases. Which EBS volume type is best for high performance NoSQL cluster deployments?

- A. io1
- B. gp1
- C. standard
- D. gp2

Answer: A

Explanation:

Provisioned IOPS SSD should be used for critical business applications that require sustained IOPS performance, or more than 10,000 IOPS or 160 MiB/s of throughput per volume

This is ideal for Large database workloads, such as:

- MongoDB
- Cassandra
- MicrosoftSQL Server
- MySQL
- PostgreSQL
- Oracle

For more information on the various CBS Volume Types, please refer to the below link:

- <http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/CBSVolumeTypes.html>

NEW QUESTION 93

You need to create a simple, holistic check for your system's general availability and uptime. Your system presents itself as an HTTP-speaking API. What is the most simple tool on AWS to achieve this with?

- A. Route53 Health Checks
- B. CloudWatch Health Checks
- C. AWS ELB Health Checks
- D. EC2 Health Checks

Answer: A

Explanation:

Amazon Route 53 health checks monitor the health and performance of your web applications, web servers, and other resources. Each health check that you create

can monitor one of the following:

- The health of a specified resource, such as a web server
- The status of an Amazon Cloud Watch alarm
- The status of other health checks

For more information on Route53 Health checks, please refer to the below link:

• <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-fa-iloover.html>

NEW QUESTION 98

You need to scale an RDS deployment. You are operating at 10% writes and 90% reads, based on your logging. How best can you scale this in a simple way?

- A. Create a second master RDS instance and peer the RDS groups.
- B. Cache all the database responses on the read side with CloudFront.
- C. Create read replicas for RDS since the load is mostly reads.
- D. Create a Multi-AZ RDS installs and route read traffic to standby.

Answer: C

Explanation:

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.

Option A is invalid because you would need to maintain the synchronization yourself with a secondary instance.

Option B is invalid because you are introducing another layer unnecessarily when you already have read replica's Option D is invalid because you only use this for Standby's

For more information on Read Replica's, please refer to the below link: <https://aws.amazon.com/rds/details/read-replicas/>

NEW QUESTION 103

You are creating an application which stores extremely sensitive financial information. All information in the system must be encrypted at rest and in transit. Which of these is a violation of this policy?

- A. ELB SSL termination.
- B. ELB Using Proxy Protocol v1.
- C. CloudFront Viewer Protocol Policy set to HTTPS redirection.
- D. Telling S3 to use AES256 on the server-side.

Answer: A

Explanation:

If you use SSL termination, your servers will always get non-secure connections and will never know whether users used a more secure channel or not. If you are using Elastic beanstalk to configure the ELB, you can use the below article to ensure end to end encryption.

<http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/configuring-https-endtoend.html>

NEW QUESTION 105

Your company uses AWS to host its resources. They have the following requirements

- 1) Record all API calls and Transitions
- 2) Help in understanding what resources are there in the account
- 3) Facility to allow auditing credentials and logins

Which services would suffice the above requirements

- A. AWS Config, CloudTrail, 1AM Credential Reports
- B. CloudTrail, 1AM Credential Reports, AWS Config
- C. CloudTrail, AWS Config, 1AM Credential Reports
- D. AWS Config, 1AM Credential Reports, CloudTrail

Answer: C

Explanation:

You can use AWS CloudTrail to get a history of AWS API calls and related events for your account. This history includes calls made with the AWS Management Console, AWS Command Line Interface, AWS SDKs, and other AWS services. For more information on Cloudtrail, please visit the below URL:

• <http://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-user-guide.html>

AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. Config continuously monitors and records your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations. With Config, you can review changes in configurations and relationships between AWS resources, dive into detailed resource configuration histories, and determine your overall compliance against the configurations specified in your internal guidelines. This enables you to simplify compliance auditing, security analysis, change management, and operational troubleshooting. For more information on the config service, please visit the below URL:

• <https://aws.amazon.com/config/>

You can generate and download a credential report that lists all users in your account and the status of their various credentials, including passwords, access keys, and MFA devices. You can get a credential report from the AWS Management Console, the AWS SDKs and Command Line Tools, or the 1AM API. For more information on Credentials Report, please visit the below URL:

• http://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_getting-report.html

NEW QUESTION 106

You need to perform ad-hoc business analytics queries on well-structured data. Data comes in constantly at a high velocity. Your business intelligence team can understand SQL.

What AWS service(s) should you look to first?

- A. Kinesis Firehose + RDS
- B. Kinesis Firehose+RedShift
- C. EMR using Hive
- D. EMR running Apache Spark

Answer: B

Explanation:

Amazon Kinesis Firehose is the easiest way to load streaming data into AWS. It can capture, transform, and load streaming data into Amazon Kinesis Analytics, Amazon S3, Amazon Redshift, and Amazon Elasticsearch Service, enabling near real-time analytics with existing business intelligence tools and dashboards you're already using today. It is a fully managed service that automatically scales to match the throughput of your data and requires no ongoing administration. It can also batch, compress, and encrypt the data before loading it, minimizing the amount of storage used at the destination and increasing security.

For more information on Kinesis firehose, please visit the below URL:

- <https://aws.amazon.com/kinesis/firehose/>

Amazon Redshift is a fully managed, petabyte-scale data warehouse service in the cloud. You can start with just a few hundred gigabytes of data and scale to a petabyte or more. This enables you to use your data to acquire new insights for your business and customers. For more information on Redshift, please visit the below URL:

<http://docs.aws.amazon.com/redshift/latest/mgmt/welcome.html>

NEW QUESTION 110

Your application consists of 10% writes and 90% reads. You currently service all requests through a Route53 Alias Record directed towards an AWS ELB, which sits in front of an EC2 Auto Scaling Group. Your system is getting very expensive when there are large traffic spikes during certain news events, during which many more people request to read similar data all at the same time. What is the simplest and cheapest way to reduce costs and scale with spikes like this?

- A. Create an S3 bucket and asynchronously replicate common requests responses into S3 object
- B. When a request comes in for a precomputed response, redirect to AWS S3.
- C. Create another ELB and Auto Scaling Group layer mounted on top of the other system, adding a tier to the system
- D. Serve most read requests out of the top layer.
- E. Create a CloudFront Distribution and direct Route53 to the Distribution
- F. Use the ELB as an Origin and specify Cache Behaviours to proxy cache requests which can be served later.
- G. Create a Memcached cluster in AWS ElastiCache
- H. Create cache logic to serve requests which can be served later from the in-memory cache for increased performance.

Answer: C

Explanation:

Use CloudFront distribution for distributing the heavy reads for your application. You can create a zone apex record to point to the CloudFront distribution.

You can control how long your objects stay in a CloudFront cache before CloudFront forwards another request to your origin. Reducing the duration allows you to serve dynamic content. Increasing the duration means your users get better performance because your objects are more likely to be served directly from the edge cache. A longer duration also reduces the load on your origin.

For more information on CloudFront object expiration, please visit the below URL:

<http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/CExpiration.html>

NEW QUESTION 113

Your application's Auto Scaling Group scales up too quickly, too much, and stays scaled when traffic decreases. What should you do to fix this?

- A. Set a longer cooldown period on the Group, so the system stops overshooting the target capacity
- B. The issue is that the scaling system doesn't allow enough time for new instances to begin servicing requests before measuring aggregate load again.
- C. Calculate the bottleneck or constraint on the compute layer, then select that as the new metric, and set the metric thresholds to the bounding values that begin to affect response latency.
- D. Raise the CloudWatch Alarms threshold associated with your autoscaling group, so the scaling takes more of an increase in demand before beginning.
- E. Use larger instances instead of lots of smaller ones, so the Group stops scaling out so much and wasting resources as the OS level, since the OS uses a higher proportion of resources on smaller instances.

Answer: B

Explanation:

The ideal case is that the right metric is not being used for the scale up and down.

Option A is not valid because it mentions that the cooldown is not happening when the traffic decreases, that means the metric threshold for the scale down is not occurring in Cloudwatch

Option C is not valid because increasing the Cloudwatch alarm metric will not ensure that the instances scale down when the traffic decreases.

Option D is not valid because the question does not mention any constraints that points to the instance size. For an example on using custom metrics for scaling in and out, please follow the below link for a use case.

- <https://blog.powerupcloud.com/aws-autoscaling-based-on-database-query-custom-metrics-f396c16e5e6a>

NEW QUESTION 116

You are planning on using encrypted snapshots in the design of your AWS Infrastructure. Which of the following statements are true with regards to EBS Encryption

- A. Snapshotting an encrypted volume makes an encrypted snapshot; restoring an encrypted snapshot creates an encrypted volume when specified / requested.
- B. Snapshotting an encrypted volume makes an encrypted snapshot when specified / requested; restoring an encrypted snapshot creates an encrypted volume when specified / requested.
- C. Snapshotting an encrypted volume makes an encrypted snapshot; restoring an encrypted snapshot always creates an encrypted volume.
- D. Snapshotting an encrypted volume makes an encrypted snapshot when specified / requested; restoring an encrypted snapshot always creates an encrypted volume.

Answer: C

Explanation:

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

Snapshots that are taken from encrypted volumes are automatically encrypted. Volumes that are created from encrypted snapshots are also automatically

encrypted.

For more information on CBS encryption, please visit the below URL:

- <http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/CBSCncryption.html>

NEW QUESTION 121

There is a very serious outage at AWS. EC2 is not affected, but your EC2 instance deployment scripts stopped working in the region with the outage. What might be the issue?

- A. The AWS Console is down, so your CLI commands do not work.
- B. S3 is unavailable, so you can't create EBS volumes from a snapshot you use to deploy new volumes.
- C. AWS turns off the DeployCode API call when there are major outages, to protect from system floods.
- D. None of the other answers make sense.
- E. If EC2 is not affected, it must be some other issue.

Answer: B

Explanation:

The CBS Snapshots are stored in S3, so if you have a script which deploys EC2 Instances, the CBS volumes need to be constructed from snapshots stored in S3.

You can back up the data on your Amazon CBS volumes to Amazon S3 by taking point-in-time snapshots. Snapshots are incremental backups, which means that only the blocks on the device that have changed after your most recent snapshot are saved. This minimizes the time required to create the snapshot and saves on storage costs by not duplicating data. When you delete a snapshot, only the data unique to that snapshot is removed. Each snapshot contains all of the information needed to restore your data (from the moment when the snapshot was taken) to a new CBS volume. For more information on CBS Snapshots, please visit the below URL:

- <http://docs.aws.amazon.com/AWSCC2/latest/UserGuide/CBSSnapshots.htm>

NEW QUESTION 124

Your company wants to understand where cost is coming from in the company's production AWS account. There are a number of applications and services running at any given time. Without expending too much initial development time, how best can you give the business a good understanding of which applications cost the most per month to operate?

- A. Create an automation script which periodically creates AWS Support tickets requesting detailed intra-month information about your bill.
- B. Use custom CloudWatch Metrics in your system, and put a metric data point whenever cost is incurred.
- C. Use AWS Cost Allocation Tagging for all resources which support it.
- D. Use the Cost Explorer to analyze costs throughout the month.
- E. Use the AWS Price API and constantly running resource inventory scripts to calculate total price based on multiplication of consumed resources over time.

Answer: C

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.

For more information on Cost Allocation tags, please visit the below URL: <http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/cost-alloc-tags.html>

NEW QUESTION 126

You need to create an audit log of all changes to customer banking data. You use DynamoDB to store this customer banking data. It's important not to lose any information due to server failures. What is an elegant way to accomplish this?

- A. Use a DynamoDB StreamSpecification and stream all changes to AWS Lambda.
- B. Log the changes to AWS CloudWatch Logs, removing sensitive information before logging.
- C. Before writing to DynamoDB, do a pre-write acknowledgment to disk on the application server, removing sensitive information before logging.
- D. Periodically rotate these log files into S3.
- E. Use a DynamoDB StreamSpecification and periodically flush to an EC2 instance store, removing sensitive information before putting the object.
- F. Periodically flush these batches to S3.
- G. Before writing to DynamoDB, do a pre-write acknowledgment to disk on the application server, removing sensitive information before logging.
- H. Periodically pipe these files into CloudWatch Logs.

Answer: A

Explanation:

You can use Lambda functions as triggers for your Amazon DynamoDB table. Triggers are custom actions you take in response to updates made to the DynamoDB table. To create a trigger, first you enable Amazon DynamoDB Streams for your table. Then, you write a Lambda function to process the updates published to the stream.

For more information on DynamoDB with Lambda, please visit the below URL: <http://docs.aws.amazon.com/lambda/latest/dg/with-ddb.html>

NEW QUESTION 128

If I want CloudFormation stack status updates to show up in a continuous delivery system in as close to real time as possible, how should I achieve this?

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

Answer: C

Explanation:

Answer - C

You can monitor the progress of a stack update by viewing the stack's events. The console's Cvents tab displays each major step in the creation and update of the stack sorted by the time of each event with latest events on top. The start of the stack update process is marked with an UPDATE_IN_PROGRCSS event for the stack. For more information on Monitoring your stack, please visit the below URL:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-updating-stacks-monitor-stack.html>

NEW QUESTION 130

What is required to achieve gigabit network throughput on EC2? You already selected cluster- compute, 10GB instances with enhanced networking, and your workload is already network-bound, but you are not seeing 10 gigabit speeds.

- A. Enable biplane networking on your servers, so packets are non-blocking in both directions and there's no switching overhead.
- B. Ensure the instances are in different VPCs so you don't saturate the Internet Gateway on any one VPC.
- C. Select PIOPS for your drives and mount several, so you can provision sufficient disk throughput.
- D. Use a placement group for your instances so the instances are physically near each other in the same Availability Zone.

Answer: D

Explanation:

A placement group is a logical grouping of instances within a single Availability Zone. Placement groups are recommended for applications that benefit from low network latency, high network throughput, or both. To provide the lowest latency, and the highest packet-per-second network performance for your placement group, choose an instance type that supports enhanced networking. For more information on Placement Groups, please visit the below URL:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>

NEW QUESTION 135

You need to deploy a new application version to production. Because the deployment is high-risk, you need to roll the new version out to users over a number of hours, to make sure everything is working correctly. You need to be able to control the proportion of users seeing the new version of the application down to the percentage point. You use ELB and EC2 with Auto Scaling Groups and custom AMIs with your code pre-installed assigned to Launch Configurations. There are no database-level changes during your deployment. You have been told you cannot spend too much money, so you must not increase the number of EC2 instances much at all during the deployment, but you also need to be able to switch back to the original version of code quickly if something goes wrong. What is the best way to meet these requirements?

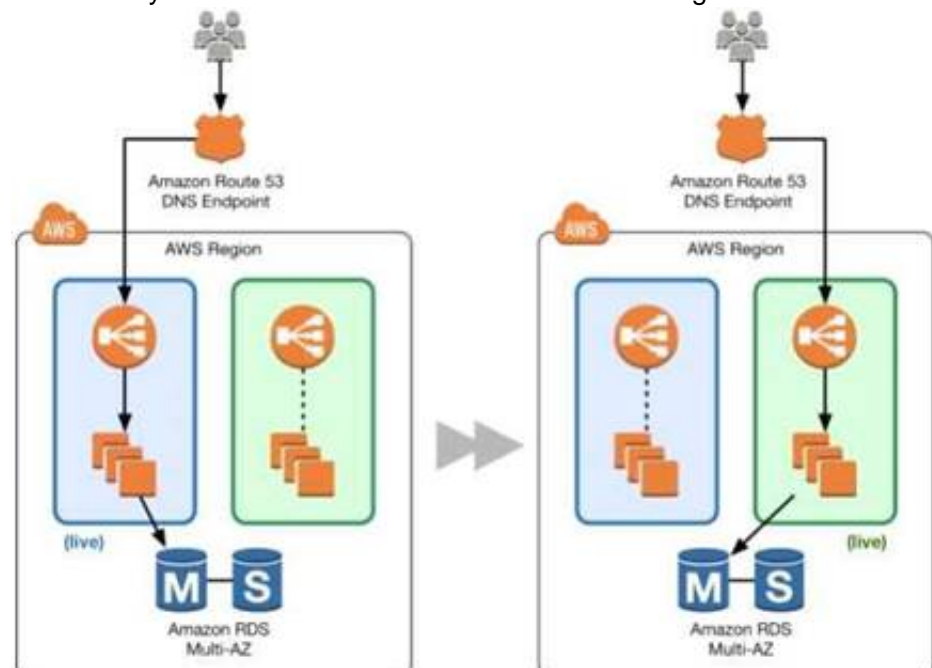
- A. Create a second ELB, Auto Scaling Launch Configuration, and Auto Scaling Group using the Launch Configuration
- B. Create AMIs with all code pre-installed
- C. Assign the new AMI to the second Auto Scaling Launch Configuration
- D. Use Route53 Weighted Round Robin Records to adjust the proportion of traffic hitting the two ELBs.
- E. Use the Blue-Green deployment method to enable the fastest possible rollback if needed
- F. Create a full second stack of instances and cut the DNS over to the new stack of instances, and change the DNS back if a rollback is needed.
- G. Create AMIs with all code pre-installed
- H. Assign the new AMI to the Auto Scaling Launch Configuration, to replace the old one
- I. Gradually terminate instances running the old code (launched with the old Launch Configuration) and allow the new AMIs to boot to adjust the traffic balance to the new code
- J. On rollback, reverse the process by doing the same thing, but changing the AMI on the Launch Config back to the original code.
- K. Migrate to use AWS Elastic Beanstalk
- L. Use the established and well-tested Rolling Deployment setting AWS provides on the new Application Environment, publishing a zip bundle of the new code and adjusting the wait period to spread the deployment over time
- M. Re-deploy the old code bundle to rollback if needed.

Answer: A

Explanation:

This is an example of a Blue Green Deployment

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're using Elastic Load Balancing



For more information on Blue Green Deployments, please visit the below URL:

- https://dOawsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf

NEW QUESTION 139

Your CTO has asked you to make sure that you know what all users of your AWS account are doing to change resources at all times. She wants a report of who is doing what over time, reported to her once per week, for as broad a resource type group as possible. How should you do this?

- A. Create a global AWS CloudTrail Trail
- B. Configure a script to aggregate the log data delivered to S3 once per week and deliver this to the CTO.
- C. Use CloudWatch Events Rules with an SNS topic subscribed to all AWS API call
- D. Subscribe the CTO to an email type delivery on this SNS Topic.
- E. Use AWS IAM credential reports to deliver a CSV of all uses of IAM UserTokens overtime to the CTO.
- F. Use AWS Config with an SNS subscription on a Lambda, and insert these changes over time into a DynamoDB table
- G. Generate reports based on the contents of this table.

Answer: A

Explanation:

AWS CloudTrail is an AWS service that helps you enable governance, compliance, and operational and risk auditing of your AWS account. Actions taken by a user, role, or an AWS service are recorded as events in CloudTrail. Events include actions taken in the AWS Management Console, AWS Command Line Interface, and AWS SDKs and APIs.

Visibility into your AWS account activity is a key aspect of security and operational best practices. You can use CloudTrail to view, search, download, archive, analyze, and respond to account activity across your AWS infrastructure. You can identify who or what took which action, what resources were acted upon, when the event occurred, and other details to help you analyze and respond to activity in your AWS account.

For more information on Cloudtrail, please visit the below URL:

- <http://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-user-guide.html>

NEW QUESTION 141

You are building a Ruby on Rails application for internal, non-production use which uses MySQL as a database. You want developers without very much AWS experience to be able to deploy new code with a single command line push. You also want to set this up as simply as possible. Which tool is ideal for this setup?

- A. AWSCloudFormation
- B. AWSOpsWorks
- C. AWS ELB+ EC2 with CLI Push
- D. AWS Elastic Beanstalk

Answer: D

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.

For more information on Elastic beanstalk, please visit the below URL:

- <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/Welcome.html>

NEW QUESTION 145

Your team wants to begin practicing continuous delivery using CloudFormation, to enable automated builds and deploys of whole, versioned stacks or stack layers. You have a 3-tier, mission-critical system. Which of the following is NOT a best practice for using CloudFormation in a continuous delivery environment?

- A. Use the AWS CloudFormation ValidateTemplate call before publishing changes to AWS.
- B. Model your stack in one template, so you can leverage CloudFormation's state management and dependency resolution to propagate all changes.
- C. Use CloudFormation to create brand new infrastructure for all stateless resources on each push, and run integration tests on that set of infrastructure.
- D. Parametrize the template and use Mappings to ensure your template works in multiple Regions.

Answer: B

Explanation:

Answer - B

Some of the best practices for Cloudformation are

- Created Nested stacks

As your infrastructure grows, common patterns can emerge in which you declare the same components in each of your templates. You can separate out these common components and create dedicated templates for them. That way, you can mix and match different templates but use nested stacks to create a single, unified stack. Nested stacks are stacks that create other stacks. To create nested stacks, use the `AWS::CloudFormation::Stackresource` in your template to reference other templates.

- Reuse Templates

After you have your stacks and resources set up, you can reuse your templates to replicate your infrastructure in multiple environments. For example, you can create environments for development, testing, and production so that you can test changes before implementing them into production. To make templates reusable, use the parameters, mappings, and conditions sections so that you can customize your stacks when you create them. For example, for your development environments, you can specify a lower-cost instance type compared to your production environment, but all other configurations and settings remain the same. For more information on Cloudformation best practises, please visit the below URL:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html>

NEW QUESTION 148

Your API requires the ability to stay online during AWS regional failures. Your API does not store any state, it only aggregates data from other sources - you do not have a database. What is a simple but effective way to achieve this uptime goal?

- A. Use a CloudFront distribution to serve up your AP
- B. Even if the region your API is in goes down, the edge locations CloudFront uses will be fine.
- C. Use an ELB and a cross-zone ELB deployment to create redundancy across datacenter
- D. Even if a region fails, the other AZ will stay online.
- E. Create a Route53 Weighted Round Robin record, and if one region goes down, have that region redirect to the other region.
- F. Create a Route53 Latency Based Routing Record with Failover and point it to two identical deployments of your stateless API in two different region
- G. Make sure both regions use Auto Scaling Groups behind ELBs.

Answer: D

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.

For more information on Route53 Failover Routing, please visit the below URL:

<http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html>

NEW QUESTION 152

You meet once per month with your operations team to review the past month's data. During the meeting, you realize that 3 weeks ago, your monitoring system which pings over HTTP from outside AWS recorded a large spike in latency on your 3-tier web service API. You use DynamoDB for the database layer, ELB, EBS, and EC2 for the business logic tier, and SQS, ELB, and EC2 for the presentation layer. Which of the following techniques will NOT help you figure out what happened?

- A. Check your CloudTrail log history around the spike's time for any API calls that caused slowness.
- B. Review CloudWatch Metrics for one minute interval graphs to determine which components slowed the system down.
- 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: B

Explanation:

The Cloudwatch metric retention is as follows. If the data points are of a one minute interval, then the graphs will not be available in Cloudwatch

- Data points with a period of less than 60 seconds are available for 3 hours. These data points are high-resolution custom metrics.
- Data points with a period of 60 seconds (1 minute) are available for 15 days
- Data points with a period of 300 seconds (5 minute) are available for 63 days
- Data points with a period of 3600 seconds (1 hour) are available for 455 days (15 months) For more information on Cloudwatch metrics, please visit the below URL:
- http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/cloudwatch_concepts.html

NEW QUESTION 153

Your CTO thinks your AWS account was hacked. What is the only way to know for certain if there was unauthorized access and what they did, assuming your hackers are very sophisticated AWS engineers and doing everything they can to cover their tracks?

- A. Use CloudTrail Log File Integrity Validation.
- B. Use AWS Config SNS Subscriptions and process events in real time.
- C. Use CloudTrail backed up to AWS S3 and Glacier.
- D. Use AWS Config Timeline forensics.

Answer: A

Explanation:

To determine whether a log file was modified, deleted, or unchanged after CloudTrail delivered it, you can use CloudTrail log file integrity validation. This feature is built using industry standard algorithms: SHA-256 for hashing and SHA-256 with RSA for digital signing. This makes it computationally infeasible to modify, delete or forge CloudTrail log files without detection. You can use the AWS CLI to validate the files in the location where CloudTrail delivered them

Validated log files are invaluable in security and forensic investigations. For example, a validated log file enables you to assert positively that the log file itself has not changed, or that particular user credentials performed specific API activity. The CloudTrail log file integrity validation process also lets you know if a log file has been deleted or changed, or assert positively that no log files were delivered to your account during a given period of time.

For more information on Cloudtrail log file validation, please visit the below URL:

<http://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html>

NEW QUESTION 154

Your serverless architecture using AWS API Gateway, AWS Lambda, and AWS DynamoDB experienced a large increase in traffic to a sustained 3000 requests per second, and dramatically increased in failure rates. Your requests, during normal operation, last 500 milliseconds on average. Your DynamoDB table did not exceed 50% of provisioned throughput, and Table primary keys are designed correctly. What is the most likely issue?

- A. Your API Gateway deployment is throttling your requests.
- B. Your AWS API Gateway Deployment is bottlenecking on request (deserialization).
- C. You did not request a limit increase on concurrent Lambda function executions.
- D. You used Consistent Read requests on DynamoDB and are experiencing semaphore lock.

Answer: C

Explanation:

Every Lambda function is allocated with a fixed amount of specific resources regardless of the memory allocation, and each function is allocated with a fixed amount of code storage per function and per account.

By default, AWS Lambda limits the total concurrent executions across all functions within a given region to 1000.

For more information on Concurrent executions, please visit the below URL: <http://docs.aws.amazon.com/lambda/latest/dg/concurrent-executions.html>

NEW QUESTION 157

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?

- A. 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 AMI based on the spot instance.
- B. 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.
- C. Purchase a Light Utilization Reserved Instance to save money on the continuous integration machine
- D. Use these credits whenever you create AMIs on instances.
- E. When the CI instance receives commits, attach a new EBS volume to the CI machine

F. Perform all setup on this EBS volume so you don't need

Answer: A

Explanation:

Amazon EC2 Spot instances allow you to bid on spare Amazon EC2 computing capacity. Since Spot instances are often available at a discount compared to On-Demand pricing, you can significantly reduce the cost of running your applications, grow your application's compute capacity and throughput for the same budget, and enable new types of cloud computing applications.

For more information on Spot Instances, please visit the below URL: <https://aws.amazon.com/ec2/spot/>

NEW QUESTION 161

Which of the following services can be used in conjunction with Cloudwatch Logs. Choose the 3 most viable services from the options given below

- A. Amazon Kinesis
- B. Amazon S3
- C. Amazon SQS
- D. Amazon Lambda

Answer: ABD

Explanation:

The AWS Documentation the following products which can be integrated with Cloudwatch logs

1) Amazon Kinesis - Here data can be fed for real time analysis

2) Amazon S3 - You can use CloudWatch Logs to store your log data in highly durable storage such as S3.

3) Amazon Lambda - Lambda functions can be designed to work with Cloudwatch log For more information on Cloudwatch Logs, please refer to the below link:
link:<http://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/WhatIsCloudWatchLogs.html>

NEW QUESTION 162

You are deciding on a deployment mechanism for your application. Which of the following deployment mechanisms provides the fastest rollback after failure.

- A. Rolling-Immutable
- B. Canary
- C. Rolling-Mutable
- D. Blue/Green

Answer: D

Explanation:

In Blue Green Deployments, you will always have the previous version of your application available.

So anytime there is an issue with a new deployment, you can just quickly switch back to the older version of your application.

For more information on Blue Green Deployments, please refer to the below link: <https://docs.cloudfoundry.org/devguide/deploy-apps/blue-green.html>

NEW QUESTION 167

There is a requirement for a vendor to have access to an S3 bucket in your account. The vendor already has an AWS account. How can you provide access to the vendor on this bucket.

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

Answer: C

Explanation:

The AWS documentation mentions

You share resources in one account with users in a different account. By setting up cross-account access in this way, you don't need to create individual IAM users in each account. In addition, users don't have to sign out of one account and sign into another in order to access resources that are in different AWS accounts. After configuring the role, you see how to use the role from the AWS Management Console, the AWS CLI, and the API

For more information on Cross Account Roles Access, please refer to the below link:

- http://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

NEW QUESTION 169

You have launched a CloudFormation template, but are receiving a failure notification after the template was launched. What is the default behavior of CloudFormation in such a case

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

Answer: A

Explanation:

The AWS Documentation mentions

AWS CloudFormation ensures all stack resources are created or deleted as appropriate. Because AWS CloudFormation treats the stack resources as a single unit,

they must all be created or deleted successfully for the stack to be created or deleted. If a resource cannot be created, AWS CloudFormation rolls the stack back and automatically deletes any resources that were created.

For more information on Cloudformation, please refer to the below link: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacks.html>

NEW QUESTION 170

Which of the following credentials types are supported by AWSCodeCommit? Select 3 Options

- A. Git Credentials
- B. SSH Keys
- C. User name/password
- D. AWS Access Keys

Answer: ABD

Explanation:

The AWS documentation mentions

I AM supports AWS CodeCommit with three types of credentials:

Git credentials, an IAM -generated user name and password pair you can use to communicate with AWS CodeCommit repositories over HTTPS.

SSH keys, a locally generated public-private key pair that you can associate with your IAM user to communicate with AWS CodeCommit repositories over SSH.

AWS access keys, which you can use with the credential helper included with the AWS CLI to communicate with AWS CodeCommit repositories over HTTPS.

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_ssh-keys.html

NEW QUESTION 175

Which of the following is the right sequence of initial steps in the deployment of application revisions using Code Deploy

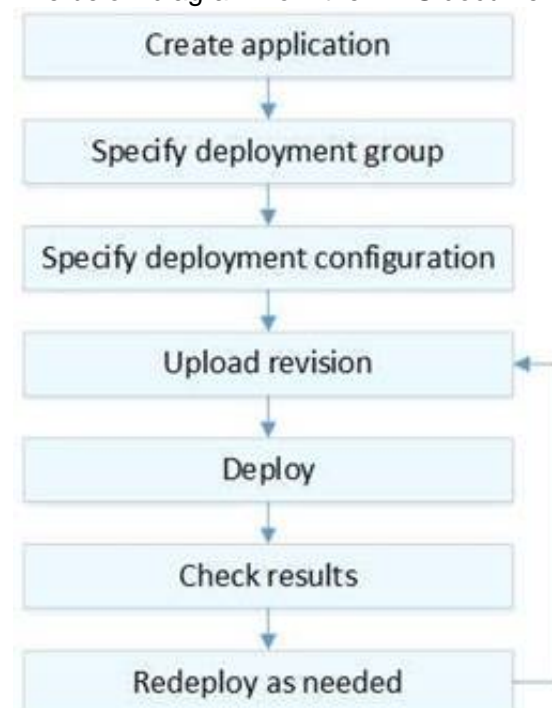
- 1) Specify deployment configuration
- 2) Upload revision
- 3) Create application
- 4) Specify deployment group

- A. 3, 2, 1 and 4
- B. 3,1,2 and 4
- C. 3,4,1 and 2
- D. 3,4,2 and 1

Answer: C

Explanation:

The below diagram from the AWS documentation shows the deployment steps



For more information on the deployment steps please refer to the below link:

- <http://docs.aws.amazon.com/codedeploy/latest/userguide/deployment-steps.html>

NEW QUESTION 177

When your application is loaded onto an Opsworks stack, which of the following event is triggered by Opsworks?

- A. Deploy
- B. Setup
- C. Configure
- D. Shutdown

Answer: A

Explanation:

When you deploy an application, AWS Ops Works Stacks triggers a Deploy event, which runs each layer's Deploy recipes. AWS OpsWorks Stacks also installs stack configuration and deployment attributes that contain all of the information needed to deploy the app, such as the app's repository and database connection data. For more information on the Deploy event please refer to the below link:

- <http://docs.aws.amazon.com/opsworks/latest/userguide/workingapps.html>

NEW QUESTION 178

Which of the following Cache Engines does Opswork have built in support for?

- A. Redis
- B. Memcache
- C. Both Redis and Memcache
- D. There is no built in support as of yet for any cache engine

Answer: B

Explanation:

The AWS Documentation mentions

AWS OpsWorks Stacks provides built-in support for Memcached. However, if Redis better suits your requirements, you can customize your stack so that your application servers use OastlCache Redis. Although it works with Redis clusters, AWS clearly specifies that AWS Opsworks stacks provide built in support for Memcached.

Amazon OastlCache is an AWS service that makes it easy to provide caching support for your application server, using either the Memcached or Redis caching engines. OastlCache can be used to improve the application server performance running on AWS Opsworks stacks.

For more information on Opswork and Cache engines please refer to the below link:

- <http://docs.aws.amazon.com/opsworks/latest/userguide/other-services-redis.html>

NEW QUESTION 181

Which of the following services can be used to implement DevOps in your company.

- A. AWS Elastic Beanstalk
- B. AWSOpswork
- C. AWS Cloudformation
- D. All of the above

Answer: D

Explanation:

All of the services can be used to implement Devops in your company

1) AWS Elastic Beanstalk, an easy-to-use service for deploying and scaling web applications and services developed with Java, .NET, PHP, Node.js, Python, Ruby, Go, and Docker on servers such as Apache, Nginx, Passenger, and IIS.

2) AWS Ops Works, a configuration management service that helps you configure and operate applications of all shapes and sizes using Chef

3) AWS Cloud Formation, which is an easy way to create and manage a collection of related AWS resources, provisioning and updating them in an orderly and predictable fashion.

For more information on AWS Devops please refer to the below link:

- <http://docs.aws.amazon.com/devops/latest/gsg/welcome.html>

NEW QUESTION 183

You have a video processing application hosted in AWS. The video's are uploaded by users onto the site. You have a program that is custom built to process those videos. The program is able to recover in case there are any failures when processing the videos. Which of the following mechanisms can be used to deploy the instances for carrying out the video processing activities, ensuring that the cost is kept at a minimum.

- A. Create a launch configuration with Reserved Instance
- B. Ensure the User Data section details the installation of the custom software
- C. Create an Autoscalinggroup with the launch configuration.
- D. Create a launch configuration with Spot Instance
- E. Ensure the User Data section details the installation of the custom software
- F. Create an Autoscalinggroupwith the launch configuration.
- G. Create a launch configuration with Dedicated Instance
- H. Ensure the User Data section details the installation of the custom software
- I. Create an Autoscaling group with the launch configuration.
- J. Create a launch configuration with On-Demand Instance
- K. Ensure the User Data section details the installation of the custom software
- L. Create an Autoscaling group with the launch configuration.

Answer: B

Explanation:

Since the application can recover from failures and cost is the priority, then Spot instances are the best bet for this requirement. The launch configuration has the facility to request for Spot Instances.

The below snapshot from the Launch configuration section shows that Spot Instances can be used for AutoScaling Groups.

Create Launch Configuration

Name	<input type="text" value="Demo"/>
Purchasing option	<input checked="" type="checkbox"/> Request Spot Instances
Current price	ap-southeast-1a 0.0173 ap-southeast-1b 0.0198
Maximum price	\$ (e.g. 0.045 = 4.5 cents/hour)
IAM role	<input type="text" value="None"/>
Monitoring	<input type="checkbox"/> Enable CloudWatch detailed monitoring Learn more
EBS-optimized instance	<input type="checkbox"/> Launch as EBS-optimized instance Additional charges apply.

► Advanced Details

For more information on Spot Instances and Autoscaling, please visit the below URL:

- <http://docs.aws.amazon.com/autoscaling/latest/userguide/US-SpotInstances.html>

NEW QUESTION 184

You currently have an Autoscalinggroup that has the following settings Min capacity-2

Desired capacity - 2 Maximum capacity - 2

Your launch configuration has AMI'S which are based on the t2.micro instance type. The application running on these instances are now experiencing issues and you have identified that the solution is to change the instance type of the instances running in the Autoscaling Group.

Which of the below solutions will meet this demand.

- Change the Instance type in the current launch configuratio
- Change the Desired value of the Autoscaling Group to 4. Ensure the new instances are launched.
- Delete the current Launch configuratio
- Create a new launch configuration with the new instance type and add it to the Autoscaling Grou
- This will then launch the new instances.
- Make a copy the Launch configuratio
- Change the instance type in the new launch configuratio
- Attach that to the Autoscaling Group.Change the maximum and Desired size of the Autoscaling Group to 4. Once the new instances are launched, change the Desired and maximum size back to 2.
- Change the desired and maximum size of the Autoscaling Group to 4. Make a copy the Launch configuratio
- Change the instance type in the new launch configuratio
- Attach that to the Autoscaling Grou
- Change the maximum and Desired size of the Autoscaling Group to 2

Answer: C

Explanation:

You should make a copy of the launch configuration, add the new instance type. The change the Autoscaling Group to include the new instance type. Then change the Desired number of the Autoscaling Group to 4 so that instances of new instance type can be launched. Once launched, change the desired size back to 2, so that Autoscaling will delete the instances with the older configuration. Note that the assumption here is that the current instances are equally distributed across multiple AZ's because Autoscaling will first use the AZRebalance process to terminate instances.

Option A is invalid because you cannot make changes to an existing Launch configuration.

Option B is invalid because if you delete the existing launch configuration, then your application will not be available. You need to ensure a smooth deployment process.

Option D is invalid because you should change the desired size to 4 after attaching the new launch configuration.

For more information on Autoscaling Suspend and Resume, please visit the below URL: <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-suspend-resume-processes.html>

NEW QUESTION 188

You currently have an Autoscalinggroup that has the following settings Min capacity-2

Desired capacity - 2 Maximum capacity - 4

The current number of instances running in the Autoscaling 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 Autoscaling Group Which of the below 2 actions can be carried out to fulfil this requirement

- Change the Maximum capacity to 2
- Change the Desired capacity to 4
- Suspend the Launch process of the Autoscaling Group
- Change the Minimum capacity to 2

Answer: AC

Explanation:

You can temporarily suspend the creation of new instances by either reducing the Maximum capacity to 2, so that the current instances running which is 2, matches the maximum limit.

Secondly you can suspend the launch process of the Autoscaling Group The AWS Documentation mentions Scaling Processes

Amazon EC2 Auto Scaling supports the following scaling processes:

Launch Adds a new CC2 instance to the group, increasing its capacity. Warning

If you suspend Launch, this disrupts other processes. For example, you can't return an instance in a standby state to service if the Launch process is suspended, because the group can't scale. For more information on Autoscaling Suspend and Resume, please visit the below URL:

? <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-suspend-resume-processes.html>

NEW QUESTION 192

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.

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

Answer: C

Explanation:

The AWS Documentation mentions

You can monitor your AWS costs by using Cloud Watch. With Cloud Watch, you can create billing alerts that notify you when your usage of your services exceeds thresholds that you define. You specify these threshold amounts when you create the billing alerts.

When your usage exceeds these amounts, AWS sends you an email notification. You can also sign up to receive notifications when AWS prices change. For more information on billing alarms, please visit the below URL:

- <http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/monitor-charges.html>

NEW QUESTION 195

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

- A. Use AWS Cloudwatch alarms to check which EC2 Instances don't match the intended instance type.
- B. Use AWS Config to create a rule to check the EC2 Instance type
- C. Use Trusted Advisor to check which EC2 Instances don't match the intended instance type.
- D. Use VPC Flow Logs to check which EC2 Instances don't match the intended instance type.

Answer: B

Explanation:

In AWS Config, you can create a rule which can be used to check if EC2 Instances follow a particular instance type. Below is a snapshot of the output of a rule to check if EC2 instances matches the type of t2micro.

The screenshot shows the AWS Config console. At the top, a rule configuration is displayed:

- Description:** Checks whether your EC2 instances are of the specified instance types.
- Trigger type:** Configuration changes
- Scope of changes:** Resources
- Resource types:** EC2 Instance
- Config rule ARN:** arn:aws:config:ap-southeast-1:213171387512:config-rule/config-rule-1exnif
- Parameters:** instanceType: t2.micro
- Overall rule status:** Last successful invocation on August 29, 2017 at 12:06:57 PM (green checkmark). Last successful evaluation on August 29, 2017 at 12:06:59 PM (green checkmark).

Below the configuration, there is a section titled "Resources evaluated". It includes a note: "Click on the ⓘ icon to view configuration details for the resource when it was last evaluated with this rule." Below this is a table with the following data:

Resource type	Config timeline ⓘ	Compliance	Last successful invocation	Last successful evaluation	Manage resource
EC2 Instance	i-0798cb2c515b3aa50	Noncompliant	August 29, 2017 11:25:10 AM	August 29, 2017 11:25:12 AM	ⓘ
EC2 Instance	i-0f2a1120bec8db3d0	Noncompliant	August 29, 2017 12:06:57 PM	August 29, 2017 12:06:59 PM	ⓘ

For more information on AWS Config, please visit the below URL:

- <https://aws.amazon.com/config/>

NEW QUESTION 196

Your company has the requirement to set up instances running as part of an Autoscaling Group. Part of the requirement is to use Lifecycle hooks to setup custom based software's and do the necessary configuration on the instances. The time required for this setup might take an hour, or might finish before the hour is up. How should you setup lifecycle hooks for the Autoscaling Group. Choose 2 ideal actions you would include as part of the lifecycle hook.

- A. Configure the lifecycle hook to record heartbeat
- B. If the hour is up, restart the timeout period.
- C. Configure the lifecycle hook to record heartbeat
- D. If the hour is up, choose to terminate the current instance and start a new one
- E. If the software installation and configuration is complete, then restart the time period.
- F. If the software installation and configuration is complete, then send a signal to complete the launch of the instance.

Answer: AD

Explanation:

The AWS Documentation provides the following information on lifecycle hooks

By default, the instance remains in a wait state for one hour, and then Auto Scaling continues the launch or terminate process (Pending: Proceed or Terminating: Proceed). If you need more time, you can restart the timeout period by recording a heartbeat. If you finish before the timeout period ends, you can complete the

lifecycle action, which continues the launch or termination process

For more information on AWS Lifecycle hooks, please visit the below URL:

- <http://docs.aws.amazon.com/autoscaling/latest/userguide/lifecycle-hooks.html>

NEW QUESTION 198

Your company is planning to setup a wordpress application. The wordpress application will connect to a MySQL 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 fulfil this requirement.

- A. Create a MySQL RDS environment with Multi-AZ feature enabled
- B. Create a MySQL RDS environment and create a Read Replica
- C. Create multiple EC2 instances in the same A
- D. Host MySQL and enable replication via scripts between the instances.
- E. Create multiple EC2 instances in separate AZ'
- F. Host MySQL and enable replication via scripts between the instances.

Answer: AD

Explanation:

One way to ensure high availability and fault tolerant environments is to ensure Instances are located across multiple availability zones. Hence if you are hosting MySQL yourself, ensure you have instances spread across multiple AZ's

The AWS Documentation mentions the following about the multi-AZ feature

Amazon RDS provides high availability and failover support for DB instances using Multi-AZ deployments. Amazon RDS uses several different technologies to provide failover support. Multi-AZ deployments for Oracle, PostgreSQL, MySQL, and MariaDB DB instances use Amazon's failover technology

For more information on AWS Multi-AZ deployments, please visit the below URL:

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZ.html>

NEW QUESTION 203

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

- A. Create a Cloudformation template with the same resources as those in the Elastic beanstalk environmen
- B. If the deployment fails, deploy the Cloudformation template.
- C. Use Rolling updates in Elastic Beanstalk so that if the deployment fails, the rolling updates feature would roll back to the last deployment.
- D. Create another parallel environment in elastic beanstal
- E. Use the Swap URL feature.
- F. Create another parallel environment in elastic beanstal
- G. Create a new Route53 Domain name for the new environment and release that url to the users.

Answer: C

Explanation:

Since the requirement is to have the least amount of downtime, the ideal way is to create a blue green deployment environment and then use the Swap URL feature

to swap environments for the new deployment and then do the swap back, incase the deployment fails.

The AWS Documentation mentions the following on the SWAP url feature of Elastic Beanstalk

Because Elastic Beanstalk performs an in-place update when you update your application versions, your application may become unavailable to users for a short period of time. It is possible to avoid this downtime by performing a blue/green deployment, where you deploy the new version to a separate environment, and then swap CNAMEs of the two environments to redirect traffic to the new version instantly.

NEW QUESTION 208

Your application is having a very high traffic, so you have enabled autoscaling in multi availability zone to suffice the needs of your application but you observe that one of the availability zone is not receiving any traffic. What can be wrong here?

- A. Autoscaling only works for single availability zone
- B. Autoscaling can be enabled for multi AZ only in north Virginia region
- C. Availability zone is not added to Elastic load balancer
- D. Instances need to manually added to availability zone

Answer: C

Explanation:

When you add an Availability Zone to your load balancer. Elastic Load Balancing creates a load balancer node in the Availability Zone. Load balancer nodes accept traffic from clients and forward requests to the healthy registered instances in one or more Availability Zones.

For more information on adding AZ's to CLB, please refer to the below U RL:

<http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/enable-disable-az.html>

NEW QUESTION 210

Your company has an on-premise Active Directory setup in place. The company has extended their footprint on AWS, but still want to have the ability to use their on-premise Active Directory for authentication. Which of the following AWS services can be used to ensure that AWS resources such as AWS Workspaces can continue to use the existing credentials stored in the on-premise Active Directory.

- A. Use the Active Directory service on AWS
- B. Use the AWS Simple AD service
- C. Use the Active Directory connector service on AWS
- D. Use the ClassicLink feature on AWS

Answer: C

Explanation:

The AWS Documentation mentions the following

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.

For more information on the AD connector, please refer to the below URL: http://docs.aws.amazon.com/directoryservice/latest/admin-guide/directory_ad_connector.html

NEW QUESTION 212

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your DOP-C01 Exam with Our Prep Materials Via below:

<https://www.certleader.com/DOP-C01-dumps.html>