

Amazon-Web-Services

Exam Questions DOP-C01

AWS Certified DevOps Engineer- Professional



NEW QUESTION 1

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

- A. You do not have the appropriate permissions to access the logs
- B. You do not have your CloudWatch metrics correctly configured
- C. ELB Access logs are only available for a maximum of one week.
- D. Access logging is an optional feature of Elastic Load Balancing that is disabled by default

Answer: D

Explanation:

Elastic Load Balancing provides access logs that capture detailed information about requests sent to your load balancer. Each log contains information such as the time the request was received, the client's IP address, latencies, request paths, and server responses. You can use these access logs to analyze traffic patterns and to troubleshoot issues.

Access logging is an optional feature of Elastic Load Balancing that is disabled by default. After you enable access logging for your load balancer, Elastic Load Balancing captures the logs and stores them in the Amazon S3 bucket that you specify. You can disable access logging at any time.

For more information on CLB access logs, please refer to the below document link: from AWS

<http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/access-log-collection.html>

NEW QUESTION 2

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 3

You have the following application to be setup in AWS

- 1) A web tier hosted on EC2 Instances
- 2) Session data to be written to DynamoDB
- 3) Log files to be written to Microsoft SQL Server

How can you allow an application to write data to a DynamoDB table?

- A. Add an IAM user to a running EC2 instance.
- B. Add an IAM user that allows write access to the DynamoDB table.
- C. Create an IAM role that allows read access to the DynamoDB table.
- D. Create an IAM role that allows write access to the DynamoDB table.

Answer: D

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. For more information on IAM Roles please refer to the below link:

<http://docs.aws.amazon.com/AWSC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

NEW QUESTION 4

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 5

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 Elastic Beanstalk, please visit the below URL:

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

NEW QUESTION 6

You use Amazon Cloud Watch as your primary monitoring system for your web application. After a recent software deployment, your users are getting Intermittent 500 Internal Server Errors when using the web application. You want to create a Cloud Watch alarm, and notify an on-call engineer when these occur. How can you accomplish this using AWS services? Choose three answers from the options given below

- A. Deploy your web application as an AWS Elastic Beanstalk applicatio
- B. Use the default Elastic Beanstalk Cloudwatch metrics to capture 500 Internal Server Error
- C. Set a CloudWatch alarm on that metric.
- D. Install a CloudWatch Logs Agent on your servers to stream web application logs to CloudWatch.
- E. Use Amazon Simple Email Service to notify an on-call engineer when a CloudWatch alarm is triggered.
- F. Create a CloudWatch Logs group and define metric filters that capture 500 Internal Server Error
- G. Set a CloudWatch alarm on that metric.
- H. Use Amazon Simple Notification Service to notify an on-call engineer when a CloudWatch alarm is triggered.

Answer: BDE

Explanation:

You can use Cloud Watch Logs to monitor applications and systems using log data

Cloud Watch 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, Cloud Watch Logs reports the data to a CloudWatch metric that you specify. Log data is encrypted while in transit and while it is at rest

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 SNS and Cloudwatch logs please refer to the below link:

http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/US_SetupSNS.html

NEW QUESTION 7

One of the instances in your Auto Scaling group health check returns the status of Impaired to Auto Scaling. What will Auto Scaling do in this case.

- A. Terminate the instance and launch a new instance
- B. Send an SNS notification
- C. Perform a health check until cool down before declaring that the instance has failed
- D. Wait for the instance to become healthy before sending traffic

Answer: A

Explanation:

Auto Scaling periodically performs health checks on the instances in your Auto Scaling group and identifies any instances that are unhealthy. You can configure Auto Scaling to determine the health status of an instance using Amazon EC2 status checks, Elastic Load Balancing health checks, or custom health checks. By default, Auto Scaling health checks use the results of the EC2 status checks to determine the health status of an instance. Auto Scaling marks an instance as unhealthy if its instance fails one or more of the status checks.

For more information monitoring in Autoscaling, please visit the below URL: <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-monitoring-features.html>

NEW QUESTION 8

You have been tasked with deploying a scalable distributed system using AWS OpsWorks. Your distributed system is required to scale on demand. As it is distributed, each node must hold a configuration file that includes the hostnames of the other instances within the layer. How should you configure AWS OpsWorks to manage scaling this application dynamically?

- A. Create a Chef Recipe to update this configuration file, configure your AWS OpsWorks stack to use custom cookbooks, and assign this recipe to the Configure Lifecycle Event of the specific layer.
- B. Update this configuration file by writing a script to poll the AWS OpsWorks service API for new instance
- C. Configure your base AMI to execute this script on Operating System startup.
- D. Create a Chef Recipe to update this configuration file, configure your AWS OpsWorks stack to use custom cookbooks, and assign this recipe to execute when instances are launched.
- E. Configure your AWS OpsWorks layer to use the AWS-provided recipe for distributed host configuration, and configure the instance hostname and file path parameters in your recipes settings.

Answer: A

Explanation:

Please check the following AWS DOCs which provides details on the scenario. Check the example of "configure".

? <https://docs.aws.amazon.com/opsworks/latest/userguide/workingcookbook-events.html> You can use the Configure Lifecycle event

This event occurs on all of the stack's instances when one of the following occurs:

- An instance enters or leaves the online state.
- You associate an Elastic IP address with an instance or disassociate one from an instance.
- You attach an Elastic Load Balancing load balancer to a layer, or detach one from a layer. Ensure the Opswork layer uses a custom Cookbook

2. Toggle **Use custom Chef cookbooks** to **Yes**.

The screenshot shows the AWS OpsWorks console configuration for a stack. The 'Use custom Chef cookbooks' toggle is set to 'Yes'. Below this, the 'Repository type' is set to 'Git', the 'Repository URL' is 'https://github.com/aws-labs/opsworks-cookbooks', and the 'Repository SSH key' is 'Optional'. The 'Branch/Revision' is also set to 'Optional'. The 'Stack color' is selected as a blue square. At the bottom, there are two diagrams of an 'AWS Region' containing an 'Amazon RDS Multi-AZ' instance, each with a blue 'M' icon and a blue 'S' icon.

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

- http://docs.aws.amazon.com/opsworks/latest/userguide/welcome_classic.html

NEW QUESTION 9

You have a set of EC2 instances hosted in AWS. You have created a role named DemoRole and assigned that role to a policy, but you are unable to use that role with an instance. Why is this the case.

- A. You need to create an instance profile and associate it with that specific role.
- B. You are not able to associate an IAM role with an instance
- C. You won't be able to use that role with an instance unless you also create a user and associate it with that specific role
- D. You won't be able to use that role with an instance unless you also create a usergroup and associate it with that specific role.

Answer: A

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.

Option B is invalid because you can associate a role with an instance

Option C and D are invalid because using users or user groups is not a pre-requisite For more information on instance profiles, please visit the link:

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

NEW QUESTION 10

You need to implement Blue/Green Deployment for several multi-tier web applications. Each of them has its individual infrastructure:

Amazon Elastic Compute Cloud (EC2) front-end servers, Amazon ElastiCache clusters, Amazon Simple Queue Service (SQS) queues, and Amazon Relational Database (RDS) Instances.

Which combination of services would give you the ability to control traffic between different deployed versions of your application?


- A. Create one AWS Elastic Beanstalk application and all AWS resources (using configuration files inside the application source bundle) for each web application
- B. New versions would be deployed using Elastic Beanstalk environments and using the Swap URLs feature.
- C. Using AWS CloudFormation templates, create one Elastic Beanstalk application and all AWS resources (in the same template) for each web application
- D. New versions would be deployed using AWS CloudFormation templates to create new Elastic Beanstalk environments, and traffic would be balanced between them using weighted Round Robin (WRR) records in Amazon Route 53. >/
- E. Using AWS CloudFormation templates, create one Elastic Beanstalk application and all AWS resources (in the same template) for each web application

- F. New versions would be deployed updating a parameter on the CloudFormation template and passing it to the cfn-hup helper daemon, and traffic would be balanced between them using Weighted Round Robin (WRR) records in Amazon Route 53.
- G. Create one Elastic Beanstalk application and all AWS resources (using configuration files inside the application source bundle) for each web application.
- H. New versions would be deployed updating the Elastic Beanstalk application version for the current Elastic Beanstalk environment.

Answer: B

Explanation:

This is an example of Blue green deployment.



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.

When it's time to promote the green environment/stack into production, update DNS records to point to the green environment/stack's load balancer. You can also do this DNS flip gradually by using the Amazon Route 53 weighted routing policy. For more information on Blue green deployment, please refer to the link:

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

NEW QUESTION 10

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.

- Use a dedicated EC2 instance with an EBS Volume.
- Download and configure the code and then create an AMI out of that.
- 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.
- Upload the code and dependencies to Amazon S3, launch an instance, download the package from Amazon S3, then create the AMI with the CreateSnapshot API call.
- 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 13

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?

- Create an Amazon S3 role in IAM with access to the specific DynamoDB tables, and assign it to the bucket hosting your website.
- Configure S3 bucket tags with your AWS access keys for your bucket hosting your website so that the application can query them for access.
- Configure a web identity federation role within IAM to enable access to the correct DynamoDB resources and retrieve temporary credentials.
- 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 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 document link: from AWS

http://docs.wsamazon.com/IAM/latest/UserGuide/id_roles_providers_oidc.html

NEW QUESTION 15

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 IAM role that allows read access to a private Amazon S3 bucket
- B. Store a file with database credentials in the Amazon S3 bucket
- 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 AMI from this instance.
- F. Store the Amazon RDS DB credentials in Amazon EC2 user data
- G. Import the credentials into the Instance on boot.
- H. Assign an IAM role to your Amazon EC2 instance, and use this IAM role to access the Amazon RDS DB from your Amazon EC2 instances.

Answer: D

Explanation:

Creating and Using an IAM Policy for IAM Database Access

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

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

The following example policy allows an IAM user to connect to a DB instance using IAM 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 IAM database authentication. They aren't valid in any other context.

IAM 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 IAM. 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.aws.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 IAM Roles, please refer to the below document link: from AWS

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

NEW QUESTION 18

You are using Chef in your data center. Which service is designed to let the customer leverage existing Chef recipes in AWS?

- A. AWS Elastic Beanstalk
- B. AWS OpsWorks
- C. AWS CloudFormation
- D. Amazon Simple Workflow Service

Answer: B

Explanation:

AWS OpsWorks is a configuration management service that uses Chef, an automation platform that treats server configurations as code. OpsWorks uses Chef to automate how servers are configured, deployed, and managed across your Amazon Elastic Compute Cloud (Amazon EC2) instances or on-premises compute environments. OpsWorks has two offerings, AWS Opsworks for Chef Automate, and AWS OpsWorks Stacks.

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

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

NEW QUESTION 20

There is a requirement to monitor API calls against your AWS account by different users and entities. There needs to be a history of those calls. The history of those calls are needed in bulk for later review. Which 2 services can be used in this scenario

- A. AWS Config; AWS Inspector
- B. AWS CloudTrail; AWS Config
- C. AWS CloudTrail; CloudWatch Events
- D. AWS Config; AWS Lambda

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>

Amazon CloudWatch Events delivers a near real-time stream of system events that describe changes in Amazon Web Services (AWS) resources. Using simple rules that you can quickly set up, you can match events and route them to one or more target functions or streams. CloudWatch Events becomes aware of operational changes as they occur. CloudWatch Events responds to these operational changes and takes corrective action as necessary, by sending messages to respond to the environment, activating functions, making changes, and capturing state information. For more information on CloudWatch events, please visit the below URL:

- <http://docs.aws.amazon.com/AmazonCloudWatch/latest/events/WhatIsCloudWatchEvents.html>

NEW QUESTION 22

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

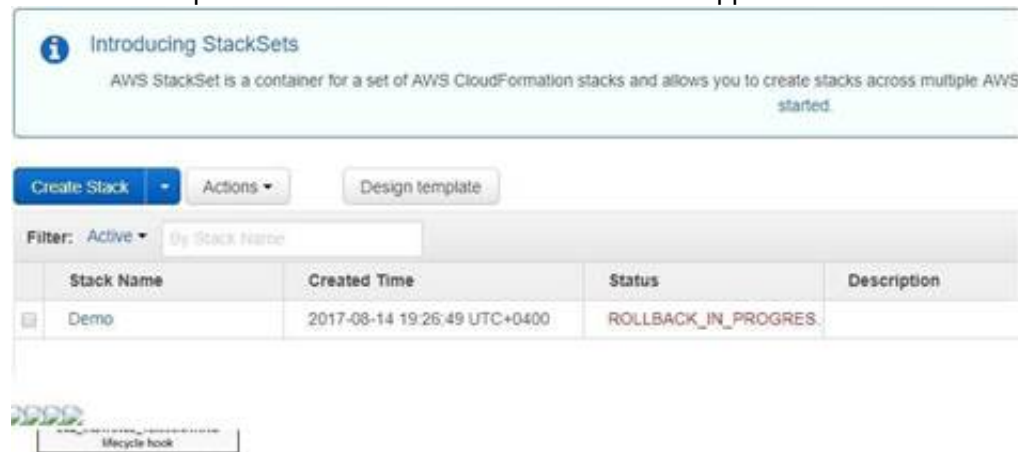
Answer: C

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 CloudFormation works, please refer to the below link: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-what-is-how-does-it-work.html>

NEW QUESTION 25

You have an asynchronous processing application using an Auto Scaling Group and an SQS Queue. The Auto Scaling Group scales according to the depth of the job queue. The completion velocity of the jobs has gone down, the Auto Scaling Group size has maxed out, but the inbound job velocity did not increase. What is a possible issue?

- A. Some of the new jobs coming in are malformed and unprocessable.
- B. The routing tables changed and none of the workers can process events anymore.
- C. Someone changed the IAM Role Policy on the instances in the worker group and broke permissions to access the queue.
- D. The scaling metric is not functioning correctly.

Answer: A

Explanation:

This question is more on the grounds of validating each option

Option B is invalid, because the Route table would have an effect on all worker processes and no jobs would have been completed.

Option C is invalid because if the IAM Role was invalid then no jobs would be completed.

Option D is invalid because the scaling is happening, it's just that the jobs are not getting completed. For more information on Scaling on Demand, please visit the below URL:

- <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-scale-based-on-demand.html>

NEW QUESTION 27

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 calls
- 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 30

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 35

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 38

Your development team is using access keys to develop an application that has access to S3 and DynamoDB. A new security policy has outlined that the credentials should not be older than 2 months, and should be rotated. How can you achieve this

- A. Use the application to rotate the keys in every 2 months via the SDK
- B. Use a script which will query the date the keys are created
- C. If older than 2 months, delete them and recreate new keys
- D. Delete the user associated with the keys after every 2 months
- E. Then recreate the user again.
- D- Delete the IAM Role associated with the keys after every 2 months
- F. Then recreate the IAM Role again.

Answer: B

Explanation:

One can use the CLI command `list-access-keys` to get the access keys. This command also returns the "CreateDate" of the keys. If the CreateDate is older than 2 months, then the keys can be deleted.

The `list-access-keys` CLI command returns information about the access key IDs associated with the specified IAM user. If there are none, the action returns

an empty list.

For more information on the CLI command, please refer to the below link: <http://docs.aws.amazon.com/cli/latest/reference/iam/list-access-keys.html>

NEW QUESTION 43

You have an application hosted in AWS, which sits on EC2 Instances behind an Elastic Load Balancer. You have added a new feature to your application and are now receiving complaints from users that the site has a slow response. Which of the below actions can you carry out to help you pinpoint the issue

- A. Use Cloudtrail to log all the API calls, and then traverse the log files to locate the issue
- B. Use Cloudwatch, monitor the CPU utilization to see the times when the CPU peaked
- C. Review the Elastic Load Balancer logs
- D. Create some custom Cloudwatch metrics which are pertinent to the key features of your application

Answer: D

Explanation:

Since the issue is occurring after the new feature has been added, it could be relevant to the new feature.
Enabling Cloudtrail will just monitor all the API calls of all services and will not benefit the cause.
The monitoring of CPU utilization will just verify that there is an issue but will not help pinpoint the issue.
The Elastic Load Balancer logs will also just verify that there is an issue but will not help pinpoint the issue.
For more information on custom Cloudwatch metrics, please refer to the below link:
<http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/publishingMetrics.html>

NEW QUESTION 44

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 49

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 54

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 59

When building a multicontainer Docker platform using Elastic Beanstalk, which of the following is required

- A. DockerFile to create custom images during deployment
- B. Prebuilt Images stored in a public or private online image repository.
- C. Kubernetes to manage the docker containers.
- D. RedHatOpenshift to manage the docker containers.

Answer: B

Explanation:

This is a special note given in the AWS Documentation for Multicontainer Docker platform for Elastic Beanstalk
Building custom images during deployment with a Dockerfile is not supported by the multicontainer Docker platform on Elastic Beanstalk. Build your images and deploy them to an online repository before creating an Elastic Beanstalk environment.
For more information on Multicontainer Docker platform for Elastic Beanstalk, please refer to the below link:
http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker_ecs.html

NEW QUESTION 60

Which of the following Deployment types are available in the CodeDeploy service. Choose 2 answers from the options given below

- A. In-place deployment
- B. Rolling deployment
- C. Immutable deployment
- D. Blue/green deployment

Answer: AD

Explanation:

The following deployment types are available

1. In-place deployment: The application on each instance in the deployment group is stopped, the latest application revision is installed, and the new version of the application is started and validated.
2. Blue/green deployment: The instances in a deployment group (the original environment) are replaced by a different set of instances (the replacement environment)

For more information on Code Deploy please refer to the below link:

- <http://docs.aws.amazon.com/codedeploy/latest/userguide/primary-components.html>

NEW QUESTION 64

You have an Opswork stack setup in AWS. You want to install some updates to the Linux instances in the stack. Which of the following can be used to publish those updates. Choose 2 answers from the options given below

- A. Create and start new instances to replace your current online instance
- B. Then delete the current instances.
- C. Use Auto-scaling to launch new instances and then delete the older instances
- D. On Linux-based instances in Chef 11.10 or older stacks, run the Update Dependencies stack command
- E. Delete the stack and create a new stack with the instances and their relevant updates

Answer: AC

Explanation:

As per AWS documentation.

By default, AWS OpsWorks Stacks automatically installs the latest updates during setup, after an instance finishes booting. AWS OpsWorks Stacks does not automatically install updates after an instance is online, to avoid interruptions such as restarting application servers. Instead, you manage updates to your online instances yourself, so you can minimize any disruptions.

We recommend that you use one of the following to update your online instances.

- Create and start new instances to replace your current online instances. Then delete the current instances.

The new instances will have the latest set of security patches installed during setup.

- On Linux-based instances in Chef 11.10 or older stacks, run the Update Dependencies stack command, which installs the current set of security patches and other updates on the specified instances.

More information is available at: <https://docs.aws.amazon.com/opsworks/latest/userguide/workingsecurity-updates.html>

NEW QUESTION 69

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.

- A. Change the Instance type in the current launch configuratio
- B. Change the Desired value of the Autoscaling Group to 4. Ensure the new instances are launched.
- C. Delete the current Launch configuratio
- D. Create a new launch configuration with the new instance type and add it to the Autoscaling Grou
- E. This will then launch the new instances.
- F. Make a copy the Launch configuratio
- G. Change the instance type in the new launch configuratio
- H. 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.
- I. Change the desired and maximum size of the Autoscaling Group to 4. Make a copy the Launch configuratio
- J. Change the instance type in the new launch configuratio
- K. Attach that to the Autoscaling Grou
- L. 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. Then 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>

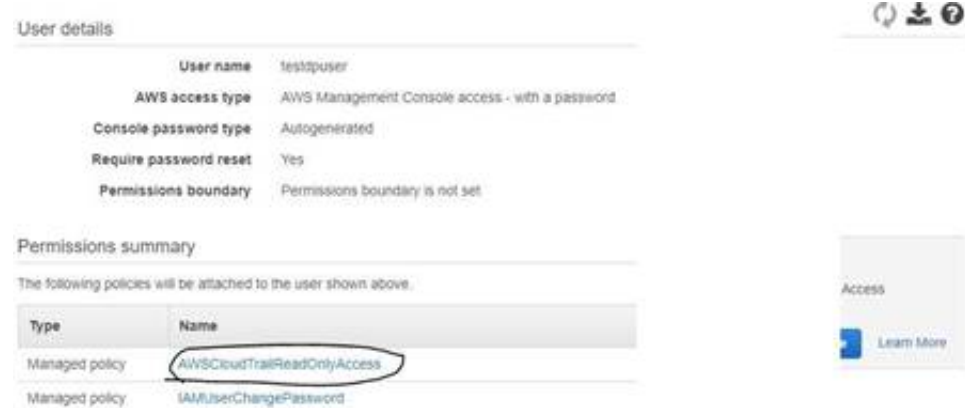
NEW QUESTION 73

An audit is going to be conducted for your company's AWS account. Which of the following steps will ensure that the auditor has the right access to the logs of your AWS account

- A. Enable S3 and ELB log
- B. Send the logs as a zip file to the IT Auditor.
- C. Ensure CloudTrail is enable
- D. Create a user account for the Auditor and attach the AWSCloudTrailReadOnlyAccess Policy to the user.
- E. Ensure that Cloudtrail is enable
- F. Create a user for the IT Auditor and ensure that full control is given to the userfor Cloudtrail.D- Enable Cloudwatch log
- G. Create a user for the IT Auditor and ensure that full control is given to the userfor the Cloudwatch logs.

Answer: B

Explanation:



The AWS Documentation clearly mentions the below

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.

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

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

NEW QUESTION 76

You are the IT administrator for your company. You have the responsibility of creating development environments which would confirm to the LAMP development stack. The requirement is that the development team always gets the latest version of the LAMP stack each time a new instance is launched. Which of the following is an efficient and effective way to implement this requirement? Choose 2 answers from the options given below

- A. Create an AMI with all the artifacts of the LAMP stack and provide an instance to the development team based on the AMI.
- B. Create a cloudformation template and use the cloud-init directives to download and the install the LAMP stack packages.
- C. Use the User data section and use a custom script which will be used to download the necessary LAMP stack packages.
- D. Create an EBS Volume with the LAMP stack and attach it to an instance whenever it is required.

Answer: BC

Explanation:

Using User data and cloud-init directives you can always ensure you download the latest version of the LAMP stack and give it to the development teams. With AMI's

you will always have the same version and will need to create an AMI everytime the version of the LAMP stack changes.

The AWS Documentation mentions

When you launch an instance in Amazon EC2, you have the option of passing user data to the instance that can be used to perform common automated configuration tasks and even run scripts after the instance starts. You can pass two types of user data to Amazon EC2: shell scripts and cloud-init directives. You can

also pass this data into the launch wizard as plain text, as a file (this is useful for launching instances using the command line tools), or as base64-encoded text (for API calls).

For more information on User data please refer to the below link: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/user-data.html>

NEW QUESTION 80

Your security officer has told you that you need to tighten up the logging of all events that occur on your AWS account. He wants to be able to access all events that occur on the account across all regions quickly and in the simplest way possible. He also wants to make sure he is the only person that has access to these events in the most secure way possible. Which of the following would be the best solution to assure his requirements are met? Choose the correct answer from the options below

- A. Use CloudTrail to log all events to one S3 bucket
- B. Make this S3 bucket only accessible by your security officer with a bucket policy that restricts access to his user only and also add MFA to the policy for a further level of security
- C. ^/
- D. Use CloudTrail to log all events to an Amazon Glacier Vault
- E. Make sure the vault access policy only grants access to the security officer's IP address.
- F. Use CloudTrail to send all API calls to CloudWatch and send an email to the security officer every time an API call is made
- G. Make sure the emails are encrypted.
- H. Use CloudTrail to log all events to a separate S3 bucket in each region as CloudTrail cannot write to a bucket in a different region
- I. Use MFA and bucket policies on all the different buckets.

Answer: A

Explanation:

AWS CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of your AWS account. With CloudTrail, you can log,

continuously monitor, and retain events related to API calls across your AWS infrastructure. CloudTrail provides a history of AWS API calls for your account, including API calls made through the AWS Management Console, AWS SDKs, command line tools, and other AWS services. This history simplifies security analysis, resource change tracking, and troubleshooting.

You can design cloudtrail to send all logs to a central S3 bucket. For more information on cloudtrail, please visit the below URL:
? <https://aws.amazon.com/cloudtrail/>

NEW QUESTION 82

You are using lifecycle hooks in your AutoScaling Group. Because there is a lifecycle hook, the instance is put in the Pending:Wait state, which means that it is not available to handle traffic yet. When the instance enters the wait state, other scaling actions are suspended. After some time, the instance state is changed to Pending:Proceed, and finally InService where the instances that are part of the Autoscaling Group can start serving up traffic. But you notice that the bootstrapping process on the instances finish much earlier, long before the state is changed to Pending:Proceed.

What can you do to ensure the instances are placed in the right state after the bootstrapping process is complete?

- A. Use the complete-lifecycle-action call to complete the lifecycle action
- B. Run this command from another EC2 Instance.
- C. Use the complete-lifecycle-action call to complete the lifecycle action
- D. Run this command from the Command line interface
- E. ^C Use the complete-lifecycle-action call to complete the lifecycle action
- F. Run this command from the Simple Notification service.
- G. Use the complete-lifecycle-action call to complete the lifecycle action
- H. Run this command from a SQS queue

Answer: B

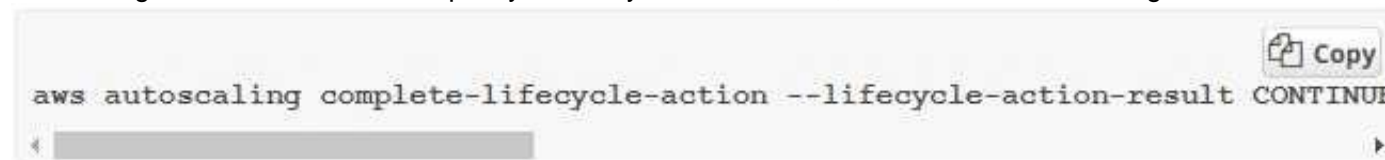
Explanation:

The AWS Documentation mentions the following

3. If you finish the custom action before the timeout period ends, use the complete-lifecycle-action command so that the Auto Scaling group can continue launching

or terminating the instance. You can specify the lifecycle action token, as shown in the following command:

3. If you finish the custom action before the timeout period ends, use the complete-lifecycle-action command so that Auto Scaling can continue launching or terminating the instance. You can specify the lifecycle action token, as shown in the following command:



For more information on lifecycle hooks, please refer to the below URL:

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

NEW QUESTION 86

You are using Autoscaling for managing the instances in your AWS environment. You need to deploy a new version of your application. You'd prefer to use all new instances if possible, but you cannot have any downtime. You also don't want to swap any environment URLs. Which of the following deployment methods would you implement

- A. Using "All at once" deployment method.
- B. Using "Blue Green" deployment method.
- C. Using "RollingUpdates" deployment method.
- D. Using "Blue Green" with "All at once" deployment method.

Answer: C

Explanation:

In Rolling deployment, you can mention a new set of servers which can replace the existing set of servers. This replacement will happen in a phased out manner. Since there is a requirement to not swap URL's, you must not use Blue Green deployments.

For more information on the differences between Rolling Updates and Blue Green deployments, please refer to the below URL:

- <https://cloudnative.io/docs/blue-green-deployment/>

NEW QUESTION 87

A custom script needs to be passed to a new Amazon Linux instances created in your Auto Scaling group. Which feature allows you to accomplish this?

- A. User data
- B. EC2Config service
- C. IAM roles
- D. AWSConfig

Answer: A

Explanation:

When you configure an instance during creation, you can add custom scripts to the User data section. So in Step 3 of creating an instance, in the Advanced Details section, we can enter custom scripts in the User Data section. The below script installs Perl during the instance creation of the EC2 instance.

Step 3: Configure Instance Details

Tenancy ⓘ Shared - Run a shared hardware instance
Additional charges will apply for dedicated tenancy.

▼ Network interfaces ⓘ

Device	Network Interface	Subnet	Primary IP	Secondary IP addresses
eni0	New network interface	subnet-95ed8dd5	Auto-assign	Add IP

Add Device

▼ Advanced Details

User data ⓘ As text ☒ As file ☐ Input is already base64 encoded

```
yum install perl-Switch perl-DateTime perl-Sys-Syslog perl-LWP-Protocol-https
-y
```

For more information on user data please refer to the URL:

- <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-metadata.html>

NEW QUESTION 89

Your IT company is currently hosting a production environment in Elastic beanstalk. You understand that the Elastic beanstalk service provides a facility known as Managed updates which are minor and patch version updates which are periodically required for your system. Your IT supervisor is worried about the impact that these updates would have on the system. What can you tell about the Elastic beanstalk service with regards to managed updates

- Package updates can be configurable weekly maintenance window
- Elastic Beanstalk applies managed updates with no downtime
- Elastic Beanstalk applies managed updates with no reduction in capacity
- All of the above

Answer: D

Explanation:

The AWS Documentation mentions the following on package updates for the Elastic beanstalk environment

You can configure your environment to apply minor and patch version updates automatically during a configurable weekly maintenance window with Managed Platform Updates. Elastic Beanstalk applies managed updates with no downtime or reduction in capacity, and cancels the update immediately if instances running your application on the new version fail health checks.

For more information on Elastic beanstalk managed updates please refer to the URL: <https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/environment-platform-update-managed.html>

<http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features/platform-upgrade.html>

NEW QUESTION 91

You were just hired as a DevOps Engineer for a startup. Your startup uses AWS for 100% of their infrastructure. They currently have no automation at all for deployment, and they have had many failures while trying to deploy to production. The company has told you deployment process risk mitigation is the most important thing now, and you have a lot of budget for tools and AWS resources.

Their stack includes a 2-tier API with data stored in DynamoDB or S3, depending on type. The Compute layer is EC2 in Auto Scaling Groups. They use Route53 for DNS pointing to an ELB. An ELB balances load across the EC2 instances. The scaling group properly varies between 4 and 12 EC2 servers. Which of the following approaches, given this company's stack and their priorities, best meets the company's needs?

- Model the stack in AWS Elastic Beanstalk as a single Application with multiple Environment
- Use Elastic Beanstalk's Rolling Deploy option to progressively roll out application code changes when promoting across environments.
- Model the stack in three CloudFormation templates: Data layer, compute layer, and networking layer
- Write stack deployment and integration testing automation following Blue-Green methodology
- >/
- Model the stack in AWS OpsWorks as a single Stack, with 1 compute layer and its associated ELB
- Use Chef and App Deployments to automate Rolling Deployment.
- Model the stack in 1 CloudFormation template, to ensure consistency and dependency graph resolution
- Write deployment and integration testing automation following Rolling Deployment methodologies.

Answer: B

Explanation:

Here you are using 2 of the best practices for deployment, one is Blue Green Deployments and the other is using Nested CloudFormation stacks.

The AWS Documentation mentions the below on 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::Stack` resource in your template to reference other templates.

For more information on CloudFormation best practices, please visit the link:

- <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html> For more information on Blue Green Deployment, please visit the link:
- https://d1awsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf

NEW QUESTION 93

Which of the following are Lifecycle events available in Opswork? Choose 3 answers from the options below

- Setup
- Decommission
- Deploy
- Shutdown

Answer: ACD

Explanation:

Below is a snapshot of the Lifecycle events in Opswork.



For more information on Lifecycle events, please refer to the below URL:

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

NEW QUESTION 95

Which of the following run command types are available for opswork stacks? Choose 3 answers from the options given below.

- A. UpdateCustom Cookbooks
- B. Execute Recipes
- C. Configure
- D. UnDeploy

Answer: ABC

NEW QUESTION 97

You work at a company that makes use of AWS resources. One of the key security policies is to ensure that all data is encrypted both at rest and in transit. Which of the following is not a right implementation which aligns to this policy?

- A. Using S3 Server Side Encryption (SSE) to store the information
- B. Enable SSL termination on the ELB
- C. Enabling Proxy Protocol
- D. Enabling sticky sessions on your load balancer

Answer: B

Explanation:

Please note the keyword "NOT" in the question.

Option A is incorrect. Enabling S3 SSE encryption helps the encryption of data at rest in S3. So Option A is invalid.

Option B is correct. If you disable SSL termination on the ELB the traffic will be encrypted all the way to the backend. SSL termination allows encrypted traffic between the client

and the ELB but cause traffic to be unencrypted between the ELB and the backend (presumably EC2 or ECS/Task, etc.)

If SSL is not terminated on the ELB you must use Layer A to have traffic encrypted all the way.

Sticky sessions are not supported with Layer A (TCP endpoint). Thus option D "Enabling sticky sessions on your load balancer" can't be used and is the right answer

For more information on sticky sessions, please visit the below URL <https://docs.aws.amazon.com/elasticloadbalancing/latest/classic/elb-sticky-sessions.html>

Requirements

- An HTTP/HTTPS load balancer.
- At least one healthy instance in each Availability Zone.
- At least one healthy instance in each Availability Zone.

If you don't want the load balancer to handle the SSL termination (known as SSL offloading), you can use TCP for both the front-end and back-end connections, and deploy certificates on the registered instances handling requests.

For more information on elb-listener-config, please visit the below

- <https://docs.aws.amazon.com/elasticloadbalancing/latest/classic/elb-listener-config.html> If the front-end connection uses TCP or SSL, then your back-end connections can use either TCP or SSL. Note: You can use an HTTPS listener and still use SSL on the backend but the ELB must terminate, decrypt and re-encrypt. This is slower and less secure than using the same encryption all the way to the backend. It also breaks the question requirement of having all data encrypted in transit since it forces the ELB to decrypt. Proxy protocol is used to provide a secure transport connection hence Option C is also incorrect. For more information on SSL Listeners for your load balancer, please visit the below URL

<http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/elb-https-load-balancers.html>

<https://aws.amazon.com/blogs/aws/elastic-load-balancer-support-for-ssl-termination/>

NEW QUESTION 98

Your company is using an Autoscaling Group to scale out and scale in instances. There is an expectation of a peak in traffic every Monday at 8am. The traffic is then expected to come down before the weekend on Friday 5pm. How should you configure Autoscaling in this?

- A. Create dynamic scaling policies to scale up on Monday and scale down on Friday
- B. Create a scheduled policy to scale up on Friday and scale down on Monday
- C. Create a scheduled policy to scale up on Monday and scale down on Friday
- D. Manually add instances to the Autoscaling Group on Monday and remove them on Friday

Answer: C

Explanation:

The AWS Documentation mentions the following for Scheduled scaling

Scaling based on a schedule allows you to scale your application in response to predictable load changes. For example, every week the traffic to your web application starts to increase on Wednesday, remains high on Thursday, and starts to decrease on Friday. You can plan your scaling activities based on the predictable traffic patterns of your web application.

For more information on scheduled scaling for Autoscaling, please visit the below URL

- http://docs.aws.amazon.com/autoscaling/latest/userguide/scheduled_scaling_time.html

NEW QUESTION 100

Explain what the following resource in a CloudFormation template does? Choose the best possible answer.

```
"SNSTopic": {  
    "Type": "AWS::SNS::Topic",  
    "Properties": {  
        "Subscription": [{  
            "Protocol": "sqs",  
            "Endpoint": { "Fn::GetAtt": [ "SQSQueue", "Arn" ] }  
        }  
    }  
}
```

- A. Creates an SNS topic which allows SQS subscription endpoints to be added as a parameter on the template
- B. Creates an SNS topic that allows SQS subscription endpoints
- C. Creates an SNS topic and then invokes the call to create an SQS queue with a logical resource name of SQSQueue
- D. Creates an SNS topic and adds a subscription ARN endpoint for the SQS resource created under the logical name SQSQueue

Answer: D

Explanation:

The intrinsic function Fn::GetAtt returns the value of an attribute from a resource in the template. This has nothing to do with adding parameters (Option A is wrong) or allowing endpoints (Option B is wrong) or invoking relevant calls (Option C is wrong)

For more information on Fn::GetAtt function please refer to the below link

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference-getatt.html>

NEW QUESTION 104

Which of the following is not a component of Elastic Beanstalk?

- A. Application
- B. Environment
- C. Docker
- D. ApplicationVersion

Answer: C

Explanation:

Answer - C

The following are the components of Elastic Beanstalk

- 1) Application - An Elastic Beanstalk application is a logical collection of Elastic Beanstalk components, including environments, versions, and environment configurations. In Elastic Beanstalk an application is conceptually similar to a folder
 - 2) Application version - In Elastic Beanstalk, an application version refers to a specific, labeled iteration of deployable code for a web application
 - 3) environment - An environment is a version that is deployed onto AWS resources. Each environment runs only a single application version at a time, however you can run the same version or different versions in many environments at the same time.
 - 4) environment Configuration - An environment configuration identifies a collection of parameters and settings that define how an environment and its associated resources behave.
 - 5) Configuration Template - A configuration template is a starting point for creating unique environment configurations. For more information on the components of Elastic Beanstalk please refer to the below link
- <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/concepts.components.html>

NEW QUESTION 109

How can you resolve a dependency Error when using CloudFormation?

- A. Use the mappings attribute
- B. Use the parameter attribute
- C. Use the DependsOn attribute
- D. Use the Error attribute

Answer: C

Explanation:

The AWS troubleshooting guide for CloudFormation states the following

To resolve a dependency error, add a DependsOn attribute to resources that depend on other resources in your template. In some cases, you must explicitly declare dependencies so that AWS CloudFormation can create or delete resources in the correct order. For example, if you create an Elastic IP and a VPC with

an Internet gateway in the same stack, the Elastic IP must depend on the Internet gateway attachment.
For more information on Cloudformation troubleshooting, please refer to the below url
<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/troubleshooting.html>

NEW QUESTION 113

You are in charge of designing a number of Cloudformation templates for your organization. You need to ensure that no one can accidentally update the production based resources on the stack during a stack update. How can this be achieved in the most efficient way?

- A. Create tags for the resources and then create IAM policies to protect the resources.
- B. Use a Stack based policy to protect the production based resources.
- C. Use S3 bucket policies to protect the resources.
- D. Use MFA to protect the resources

Answer: B

Explanation:

The AWS Documentation mentions

When you create a stack, all update actions are allowed on all resources. By default, anyone with stack update permissions can update all of the resources in the stack. During an update, some resources might require an interruption or be completely replaced, resulting in new physical IDs or completely new storage. You can prevent stack resources from being unintentionally updated or deleted during a stack update by using a stack policy. A stack policy is a JSON document that defines the update action that can be performed on designated resources.

For more information on protecting stack resources, please visit the below url <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/protect-stack-resources.html>

NEW QUESTION 118

Which of the following are true with regard to Opsworks stack Instances? Choose 3 answers from the options given below.

- A. A stack's instances can be a combination of both Linux and Windows based operating systems.
- B. You can use EC2 Instances that were created outside the boundary of Opswork.
- C. You can use instances running on your own hardware.
- D. You can start and stop instances manually.

Answer: BCD

Explanation:

The AWS Documentation mentions the following

1) You can start and stop instances manually or have AWS Ops Works Stacks automatically scale the number of instances. You can use time-based automatic scaling with any stack; Linux stacks also can use load-based scaling.

2) In addition to using AWS OpsWorks Stacks to create Amazon Linux instances, you can also register instances with a Linux stack that were created outside of AWS Ops Works Stacks. This includes Amazon EC2 instances and instances running on your own hardware. However, they must be running one of the supported Linux distributions. You cannot register Amazon EC2 or on-premises Windows instances.

3) A stack's instances can run either Linux or Windows. A stack can have different Linux versions or distributions on different instances, but you cannot mix Linux and Windows instances.

For more information on Opswork instances, please visit the below url <http://docs.aws.amazon.com/opsworks/latest/userguide/workinginstances-os.html>

NEW QUESTION 119

You want to use Code Deploy to deploy code that is hosted on your GitHub repository. Which of the following additional services can help fulfill this requirement.

- A. Use the CodePipeline service
- B. Use the CodeCommit service
- C. Use the CodeBatch service
- D. Use the SQS service

Answer: A

Explanation:

The AWS Documentation mentions the following

AWS CodePipeline is a continuous delivery service you can use to model, visualize, and automate the steps required to release your software. You can quickly model and configure the different stages of a software release process. AWS CodePipeline automates the steps required to release your software changes continuously.

For more information on CodePipeline, please visit the below URL: <http://docs.aws.amazon.com/codepipeline/latest/userguide/concepts.html>

NEW QUESTION 122

Which of the following tools for EC2 can be used to administer instances without the need to SSH or RDP into the instance.

- A. AWS Config
- B. AWS CodePipeline
- C. Run Command
- D. EC2 Config

Answer: C

Explanation:

You can use Run Command from the Amazon Linux console to configure instances without having to login to each instance

For more information on the Run Command, please visit the below URL:

- <http://docs.aws.amazon.com/systems-manager/latest/userguide/rc-console.html>

NEW QUESTION 124

An application is currently writing a large number of records to a DynamoDB table in one region. There is a requirement for a secondary application to just take in the changes to the DynamoDB table every 2 hours and process the updates accordingly. Which of the following is an ideal way to ensure the secondary application can get the relevant changes from the DynamoDB table.

- A. Insert a timestamp for each record and then scan the entire table for the timestamp as per the last 2 hours.
- B. Create another DynamoDB table with the records modified in the last 2 hours.
- C. Use DynamoDB streams to monitor the changes in the DynamoDB table.
- D. Transfer the records to S3 which were modified in the last 2 hours

Answer: C

Explanation:

The AWS Documentation mentions the following

A DynamoDB stream is an ordered flow of information about changes to items in an Amazon DynamoDB table. When you enable a stream on a table, DynamoDB captures information about every modification to data items in the table.

Whenever an application creates, updates, or deletes items in the table, DynamoDB Streams writes a stream record with the primary key attribute(s) of the items that were modified. A stream record contains information about a data modification to a single item in a DynamoDB table. You can configure the stream so that the stream records capture additional information, such as the "before" and "after" images of modified items.

For more information on DynamoDB streams, please visit the below URL: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.html>

NEW QUESTION 126

You are creating a CloudFormation template which takes in a database password as a parameter. How can you ensure that the password is not visible when anybody tries to describe the stack

- A. Use the password attribute for the resource
- B. Use the NoEcho property for the parameter value
- C. Use the hidden property for the parameter value
- D. Set the hidden attribute for the CloudFormation resource.

Answer: B

Explanation:

The AWS Documentation mentions

For sensitive parameter values (such as passwords), set the NoEcho property to true. That way, whenever anyone describes your stack, the parameter value is shown as asterisks (*•*).

For more information on CloudFormation parameters, please visit the below URL:

- <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/parameters-section-structure.html>

NEW QUESTION 131

You are creating a CloudFormation template in which UserData is going to be passed to underlying EC2 Instance. Which of the below functions is normally used to pass data to the UserData section in the CloudFormation template?

- A. "UserData": { "Fn::Base64": {
- B. "UserData": < "Fn::Ref": {
- C. "UserData": { "Fn::GetAtt": {
- D. "UserData": { "Fn::FindInMap": {

Answer: A

Explanation:

The AWS Documentation mentions

The intrinsic function Fn::Base64 returns the Base64 representation of the input string. This function is typically used to pass encoded data to Amazon EC2 instances by way of the User Data property.

For more information on the Fn::Base64 function, please visit the below URL: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference-base64.html>

NEW QUESTION 134

You have just developed a new mobile application that handles analytics workloads on large scale datasets that are stored on Amazon Redshift. Consequently, the application needs to access Amazon Redshift tables. Which of the below methods would be the best, both practically and security-wise, to access the tables?

Choose the correct answer from the options below

- A. Create an IAM user and generate encryption keys for that use
- B. Create a policy for RedShift read-only access
- C. Embed the keys in the application.
- D. Create an HSM client certificate in Redshift and authenticate using this certificate.
- E. Create a RedShift read-only access policy in IAM and embed those credentials in the application.
- F. Use roles that allow a web identity federated user to assume a role that allows access to the RedShift table by providing temporary credentials.

Answer: D

Explanation:

For access to any AWS service, the ideal approach for any application is to use Roles. This is the first preference. Hence option A and C are wrong.

For more information on IAM policies please refer to the below link: http://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies.html

Next for any web application, you need to use web identity federation. Hence option D is the right option. This along with the usage of roles is highly stressed in the AWS documentation.

"When you write such an app, you'll make requests to AWS services that must be signed with an AWS access key. However, we strongly recommend that you do not embed or distribute long-term AWS credentials with apps that a user downloads to a device, even in an encrypted store. Instead, build your app so that it requests temporary AWS security credentials dynamically when needed using web identity federation. The supplied temporary credentials map to an AWS role that has only

the permissions needed to perform the tasks required by the mobile app".

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 137

The company you work for has a huge amount of infrastructure built on AWS. However there has been some concerns recently about the security of this infrastructure, and an external auditor has been given the task of running a thorough check of all of your company's AWS assets. The auditor will be in the USA while your company's infrastructure resides in the Asia Pacific (Sydney) region on AWS. Initially, he needs to check all of your VPC assets, specifically, security groups and NACLs. You have been assigned the task of providing the auditor with a login to be able to do this. Which of the following would be the best and most secure solution to provide the auditor with so he can begin his initial investigations? Choose the correct answer from the options below

- A. Create an IAM user tied to an administrator role
- B. Also provide an additional level of security with MFA.
- C. Give him root access to your AWS Infrastructure, because he is an auditor he will need access to every service.
- D. Create an IAM user who will have read-only access to your AWS VPC infrastructure and provide the auditor with those credentials.
- E. Create an IAM user with full VPC access but set a condition that will not allow him to modify anything if the request is from any IP other than his own.

Answer: C

Explanation:

Generally you should refrain from giving high level permissions and give only the required permissions. In this case option C fits well by just providing the relevant access which is required.

For more information on IAM please see the below link:

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

NEW QUESTION 141

What are the benefits when you implement a Blue Green deployment for your infrastructure or application level changes. Choose 3 answers from the options given below

- A. Near zero-downtime release for new changes
- B. Better rollback capabilities
- C. Ability to deploy with higher risk
- D. Good turnaround time for application deployments

Answer: ABD

Explanation:

The AWS Documentation mentions the following

Blue/green deployments provide near zero-downtime release and rollback capabilities. The fundamental idea behind blue/green deployment is to shift traffic between two identical environments that are running different versions of your application. The blue environment represents the current application version serving production traffic. In parallel, the green environment is staged running a different version of your application. After the green environment is ready and tested, production traffic is redirected from blue to green.

For more information on Blue Green deployments please see the below link:

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

NEW QUESTION 142

Your team is responsible for an AWS Elastic Beanstalk application. The business requires that you move to a continuous deployment model, releasing updates to the application multiple times per day with zero downtime. What should you do to enable this and still be able to roll back almost immediately in an emergency to the previous version?

- A. Enable rolling updates in the Elastic Beanstalk environment, setting an appropriate pause time for application startup.
- B. Create a second Elastic Beanstalk environment running the new application version, and swap the environment CNAMEs.
- C. Develop the application to poll for a new application version in your code repository; download and install to each running Elastic Beanstalk instance.
- D. Create a second Elastic Beanstalk environment with the new application version, and configure the old environment to redirect clients, using the HTTP 301 response code, to the new environment

Answer: B

Explanation:

The AWS Documentation mentions the below

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. For more information on Elastic beanstalk swap URL please see the below link:

- <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.CNAMCSwap.html>

NEW QUESTION 145

You are managing the development of an application that uses DynamoDB to store JSON data. You have already set the Read and Write capacity of the DynamoDB table. You are unsure of the amount of the traffic that will be received by the application during the deployment time. How can you ensure that the DynamoDB is not highly throttled and does not become a bottleneck for the application? Choose 2 answers from the options below.

- A. Monitor the ConsumedReadCapacityUnits and ConsumedWriteCapacityUnits metric using Cloudwatch.
- B. Monitor the SystemErrors metric using Cloudwatch
- C. Create a Cloudwatch alarm which would then send a trigger to AWS Lambda to increase the Read and Write capacity of the DynamoDB table.
- D. Create a Cloudwatch alarm which would then send a trigger to AWS Lambda to create a new DynamoDB table.

Answer: AC

Explanation:

Refer to the following AWS Documentation that specifies what should be monitored for a DynamoDB table.

How can I determine how much of my provisioned throughput is being used?	You can monitor <i>ConsumedReadCapacityUnits</i> or <i>ConsumedWriteCapacityUnits</i> over the specified time period, to track how much of your provisioned throughput is being used.
--	---

For more information on monitoring DynamoDB please see the below link:

- <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/monitoring-cloudwatch.html>

NEW QUESTION 147

You are a Devops Engineer for your company. You are in charge of an application that uses EC2, ELB and Autoscaling. You have been requested to get the ELB access logs. When you try to access the logs, you can see that nothing has been recorded in S3. Why is this the case?

- A. You don't have the necessary access to the logs generated by ELB.
- B. By default ELB access logs are disabled.
- C. The Autoscaling service is not sending the required logs to ELB
- D. The EC2 Instances are not sending the required logs to ELB

Answer: B

Explanation:

The AWS Documentation mentions

Access logging is an optional feature of Elastic Load Balancing that is disabled by default. After you enable access logging for your load balancer, Elastic Load Balancing captures the logs and stores them in the Amazon S3 bucket that you specify. You can disable access logging at any time.

For more information on ELB access logs please see the below link:

- <http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/access-log-collection.html>

NEW QUESTION 151

When you implement a lifecycle hook in Autoscaling, by default what is the time limit in which the instance will be in a pending state.

- A. 60 seconds
- B. 5 minutes
- C. 60 minutes
- D. 120 minutes

Answer: C

Explanation:

The AWS Documentation mentions

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 Autoscaling lifecycle hooks please see the below link:

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

NEW QUESTION 156

Which of the following is a reliable and durable logging solution to track changes made to your AWS resources?

- A. Create a new CloudTrail trail with one new S3 bucket to store the logs and with the global services option selected
- B. Use IAM roles, S3 bucket policies and Multi-Factor Authentication (MFA) Delete on the S3 bucket that stores your logs
- C. V
- D. Create a new CloudTrail trail with one new S3 bucket to store the logs
- E. Configure SNS to send log file delivery notifications to your management system
- F. Use IAM roles and S3 bucket policies on the S3 bucket that stores your logs.
- G. Create a new CloudTrail trail with an existing S3 bucket to store the logs and with the global services option selected
- H. Use S3 ACLs and Multi-Factor Authentication (MFA) Delete on the S3 bucket that stores your logs.
- I. Create three new CloudTrail trails with three new S3 buckets to store the logs one for the AWS Management console, one for AWS SDKs and one for command line tools. Use IAM roles and S3 bucket policies on the S3 buckets that store your logs.

Answer: A

Explanation:

AWS Identity and Access Management (IAM) is integrated with AWS CloudTrail, a service that logs AWS events made by or on behalf of your AWS account. CloudTrail logs authenticated AWS API calls and also AWS sign-in events, and collects this event information in files that are delivered to Amazon S3 buckets. You need to ensure that all services are included. Hence option B is partially correct.

Option B and D is wrong because it just adds an overhead for having 3 S3 buckets and SNS notifications.

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

- <http://docs.aws.amazon.com/IAM/latest/UserGuide/cloudtrail-integration.html>

NEW QUESTION 160

Which of the following is not a supported platform for the Elastic Beanstalk service?

- A. Java
- B. AngularJS
- C. PHP
- D. .NET

Answer: B

Explanation:

Below are the supported platforms for Elastic beanstalk

Platforms

- Packer Builder
- Single Container Docker
- Multicontainer Docker
- Preconfigured Docker
- Go
- Java SE
- Java with Tomcat
- .NET on Windows Server with IIS
- Node.js
- PHP
- Python
- Ruby

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

<http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/concepts.platforms.html>

NEW QUESTION 162

You are a Devops engineer for your company. There is a requirement to host a custom application which has custom dependencies for a development team. This needs to be done using AWS service. Which of the following is the ideal way to fulfil this requirement.

- A. Package the application and dependencies with Docker, and deploy the Docker container with CloudFormation.
- B. Package the application and dependencies with Docker, and deploy the Docker container with Elastic Beanstalk.
- C. Package the application and dependencies in an S3 file, and deploy the Docker container with Elastic Beanstalk.
- D. Package the application and dependencies with in Elastic Beanstalk, and deploy with Elastic Beanstalk

Answer: B

Explanation:

The AWS Documentation mentions

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.

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

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

NEW QUESTION 163

You have a web application running on six Amazon EC2 instances, consuming about 45% of resources on each instance. You are using auto-scaling to make sure that six instances are running at all times. The number of requests this application processes is consistent and does not experience spikes. The application is critical to your business and you want high availability at all times. You want the load to be distributed evenly between all instances. You also want to use the same Amazon Machine Image (AMI) for all instances. Which of the following architectural choices should you make?

- A. Deploy 6 EC2 instances in one availability zone and use Amazon Elastic Load Balancer.
- B. Deploy 3 EC2 instances in one region and 3 in another region and use Amazon Elastic Load Balancer.
- C. Deploy 3 EC2 instances in one availability zone and 3 in another availability zone and use Amazon Elastic Load Balancer.
- D. Deploy 2 EC2 instances in three regions and use Amazon Elastic Load Balancer.

Answer: C

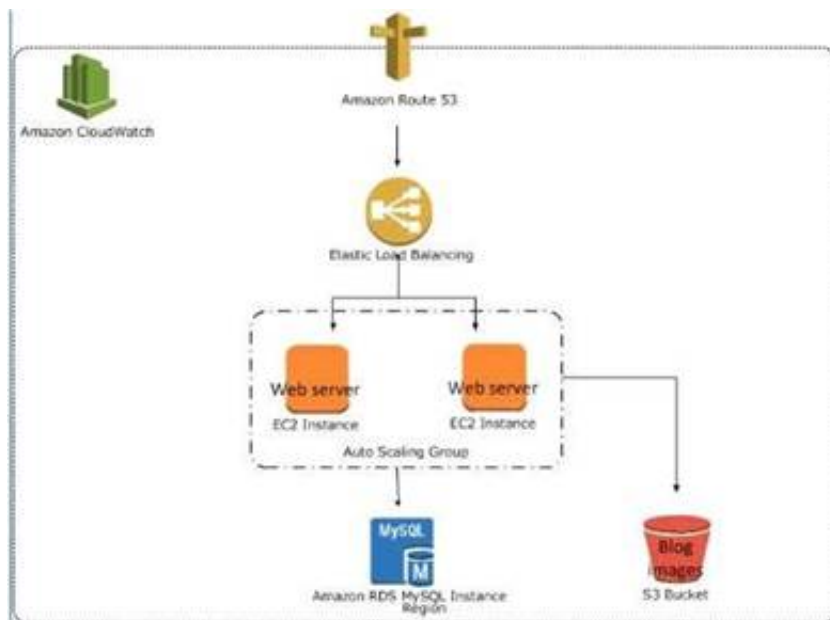
Explanation:

Option A is automatically incorrect because remember that the question asks for high availability. For option A, if the AZ goes down then the entire application fails. For Option B and D, the CLB is designed to only run in one region in AWS and not across multiple regions. So these options are wrong.

The right option is C.

The below example shows an Elastic Loadbalancer connected to 2 EC2 instances connected via Auto Scaling. This is an example of an elastic and scalable web tier.

By scalable we mean that the Auto scaling process will increase or decrease the number of EC2 instances as required.



For more information on best practices for AWS Cloud applications, please visit the below URL:

- https://d03wsstatic.com/whitepapers/AWS_Cloud_Best_Practices.pdf

NEW QUESTION 166

You are designing an application that contains protected health information. Security and compliance requirements for your application mandate that all protected health information in the application use encryption at rest and in transit. The application uses a three-tier architecture where data flows through the load balancer and is stored on Amazon EBS volumes for processing and the results are stored in Amazon S3 using the AWS SDK.

Which of the following two options satisfy the security requirements? (Select two)

- A. UseSSL termination on the load balancer, Amazon EBS encryption on Amazon EC2instances and Amazon S3 with server- side encryption.
- B. UseSSL termination with a SAN SSL certificate on the load balance
- C. Amazon EC2with all Amazon EBS volumes using Amazon EBS encryption, and Amazon S3 withserver-side encryption with customer-managed keys.
- D. UseTCP load balancing on the load balance
- E. SSL termination on the Amazon EC2instance
- F. OS- level disk encryption on the Amazon EBS volumes and Amazon S3with server-side encryption.
- G. UseTCP load balancing on the load balance
- H. SSL termination on the Amazon EC2instances and Amazon S3 with server-side encryption.
- I. UseSSL termination on the load balancer an SSL listener on the Amazon EC2instances, Amazon EBS encryption on EBS volumes containing PHI and Amazon S3with server-side encryption.

Answer: CE

Explanation:

The AWS Documentation mentions the following: HTTPS/SSL Listeners

You can create a load balancer with the following security features. SSL Server Certificates

If you use HTTPS or SSL for your front-end connections, you must deploy an X.509 certificate (SSL server certificate) on your load balancer. The load balancer decrypts

requests from clients before sending them to the back-end instances (known as SSL termination). For more information, see SSL/TLS Certificates for Classic Load Balancers.

If you don't want the load balancer to handle the SSL termination (known as SSL offloading), you can use TCP for both the front-end and back-end connections, and deploy certificates on the registered instances handling requests.

Reference Link:

? <http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/elb-listener-config.htm>

Create a Classic Load Balancer with an HTTPS Listener

A load balancer takes requests from clients and distributes them across the EC2 instances that are registered with the load balancer.

You can create a load balancer that listens on both the HTTP (80) and HTTPS (443) ports. If you specify that the HTTPS listener sends requests to the instances on port 80, the load balancer terminates the requests and communication from the load balancer to the instances is not encrypted. If the HTTPS listener sends requests to the instances on port 443, communication from the load balancer to the instances is encrypted.

Reference Link:

• <http://docs.aws.amazon.com/elasticloadbalancing/latest/classic/elb-create-https-ssl-load-balancer.htm> | Option A & B are incorrect because they are missing encryption in transit between ELB and EC2 instances.

Option D is incorrect because it is missing encryption at rest on the data associated with the EC2 instances.

NEW QUESTION 169

You are responsible for an application that leverages the Amazon SDK and Amazon EC2 roles for storing and retrieving data from Amazon S3, accessing multiple DynamoDB tables, and exchanging message with Amazon SQS queues. Your VP of Compliance is concerned that you are not following security best practices for securing all of this access. He has asked you to verify that the application's AWS access keys are not older than six months and to provide control evidence that these keys will be rotated a minimum of once every six months.

Which option will provide your VP with the requested information?

- A. Createa script to query the IAM list-access keys API to get your application accesskey creation date and create a batch process to periodically create acompliance report for your VP.
- B. Provideyour VP with a link to IAM AWS documentation to address the VP's key rotationconcerns.
- C. Updateyour application to log changes to its AWS access key credential file and use aperiodic Amazon EMR job to create a compliance report for your VP
- D. Createa new set of instructions for your configuration management tool that willperiodically create and rotate the application's existing access keys andprovide a compliance report to your VP.

Answer: B

Explanation:

The question is focusing on IAM roles rather than using access keys for accessing the services, AWS will take care of the temporary credentials provided through the roles in accessing these services.

NEW QUESTION 170

You have an I/O and network-intensive application running on multiple Amazon EC2 instances that cannot handle a large ongoing increase in traffic. The Amazon EC2 instances are using two Amazon EBS PIOPS volumes each, and each instance is identical.

Which of the following approaches should be taken in order to reduce load on the instances with the least disruption to the application?

- A. Create an AMI from each instance, and set up Auto Scaling groups with a larger instance type that has enhanced networking enabled and is Amazon EBS-optimized.
- B. Stop each instance and change each instance to a larger Amazon EC2 instance type that has enhanced networking enabled and is Amazon EBS-optimized.
- C. Ensure that RAID striping is also set up on each instance.
- D. Add an instance-store volume for each running Amazon EC2 instance and implement RAID striping to improve I/O performance.
- E. Add an Amazon EBS volume for each running Amazon EC2 instance and implement RAID striping to improve I/O performance.
- F. Create an AMI from an instance, and set up an Auto Scaling group with an instance type that has enhanced networking enabled and is Amazon EBS-optimized.

Answer: E

Explanation:

The AWS Documentation mentions the following on AMIs:

An Amazon Machine Image (AMI) provides the information required to launch an instance, which is a virtual server in the cloud. You specify an AMI when you launch

an instance, and you can launch as many instances from the AMI as you need. You can also launch instances from as many different AMIs as you need.

For more information on AMIs, please visit the link:

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

NEW QUESTION 174

Which of the following services allows you to easily run and manage Docker-enabled applications across a cluster of Amazon EC2 instances?

- A. Elastic Beanstalk
- B. Elastic Container Service
- C. Opsworks
- D. Cloudwatch

Answer: B

Explanation:

The AWS documentation provides the following information:

Amazon EC2 Container Service (ECS) allows you to easily run and manage Docker-enabled applications across a cluster of Amazon EC2 instances. Applications packaged as containers locally will deploy and run in the same way as containers managed by Amazon ECS. Amazon ECS eliminates the need to install, operate, and scale your own cluster management infrastructure, and allows you to schedule Docker-enabled applications across your cluster based on your resource needs and availability requirements.

For more information on ECS, please visit the link:

- <https://aws.amazon.com/ecs/details/>

NEW QUESTION 177

You have just been assigned to take care of the Automated resources which have been setup by your company in AWS. You are looking at integrating some of the company's chef recipes to be used for the existing Opsworks stacks already setup in AWS. But when you go to the recipes section, you cannot see the option to add any recipes. What could be the reason for this?

- A. Once you create a stack, you cannot assign custom recipe's, this needs to be done when the stack is created.
- B. Once you create layers in the stack, you cannot assign custom recipe's, this needs to be done when the layers are created.
- C. The stack layers were created without the custom cookbooks option.
- D. Just change the layer settings accordingly.
- E. The stacks were created without the custom cookbooks option.
- F. Just change the stack settings accordingly.

Answer: D

Explanation:

The AWS Documentation mentions the below:

To have a stack install and use custom cookbooks, you must configure the stack to enable custom cookbooks, if it is not already configured. You must then provide the repository URL and any related information such as a password.

For more information on Custom cookbooks for Opsworks, please visit the below URL:

- <http://docs.aws.amazon.com/opsworks/latest/userguide/workingcookbook-installingcustom-enable.html>

NEW QUESTION 182

You are a DevOps Engineer and are designing an Opsworks stack in AWS. The company has some custom recipes that are part of their on-premise Chef configuration. These same recipes need to be run whenever an instance is launched in Opsworks. Which of the following steps need to be carried out to ensure this requirement gets fulfilled. Choose 2 answers from the options given below.

- A. Ensure the custom cookbooks option is set in Opsworks stack.
- B. Ensure the custom cookbooks option is set in Opsworks layer.
- C. Ensure the recipe is placed as part of the Setup Lifecycle event as part of the Layer setting.
- D. Ensure the recipe is placed as part of the Setup Lifecycle event as part of the Stack setting.

Answer: AC

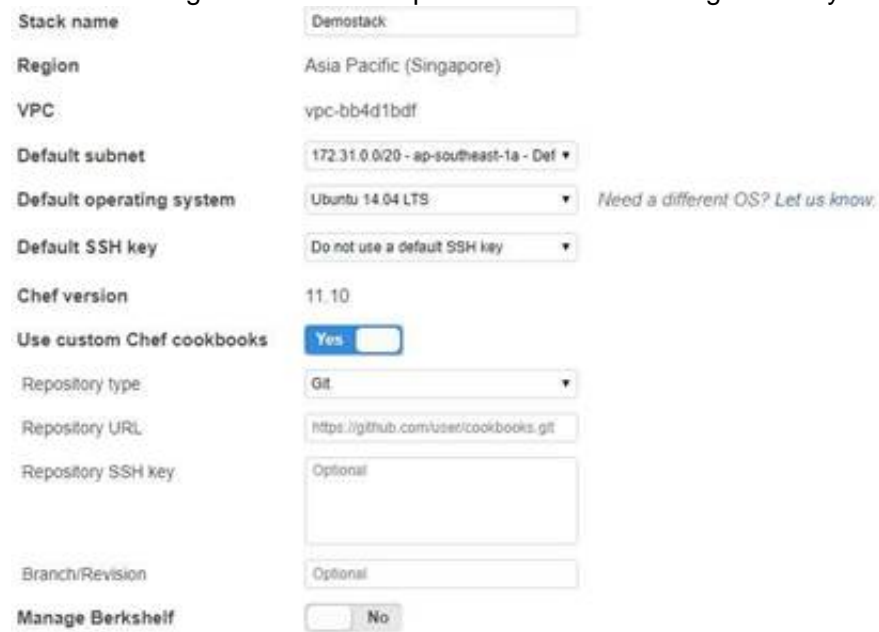
Explanation:

The AWS Documentation mentions the below:

Each layer has a set of built-in recipes assigned to each lifecycle event, although some layers lack Undeploy recipes. When a lifecycle event occurs on an instance,

AWS OpsWorks Stacks runs the appropriate set of recipes for the associated layer.

The below diagram shows a snapshot in the stack settings where you can enable custom cookbooks



The screenshot shows the 'Stack settings' for a stack named 'Demostack'. The 'Use custom Chef cookbooks' section is highlighted with a red box. The settings are as follows:

- Stack name: Demostack
- Region: Asia Pacific (Singapore)
- VPC: vpc-bb4d1bdf
- Default subnet: 172.31.0.0/20 - ap-southeast-1a - Def ▼
- Default operating system: Ubuntu 14.04 LTS ▼ *Need a different OS? Let us know.*
- Default SSH key: Do not use a default SSH key ▼
- Chef version: 11.10
- Use custom Chef cookbooks: ☒ Yes
- Repository type: Git ▼
- Repository URL: <https://github.com/user/cookbooks.git>
- Repository SSH key: Optional
- Branch/Revision: Optional
- Manage Berkshelf: ☐ No

For more information on automating recipe's, please visit the below URL: <http://docs.aws.amazon.com/opsworks/latest/userguide/workingcookbook-assigningcustom.html>

NEW QUESTION 183

Which of the following commands for the elastic beanstalk CLI can be used to create the current application into the specified environment?

- A. ebcreate
- B. ebstart
- C. enenv
- D. enapp

Answer: A

Explanation:

Differences from Version 3 of EB CLI

CB is a command line interface (CLI) tool for Clastic Beanstalk that you can use to deploy applications quickly and more easily. The latest version of CB was introduced by Clastic Beanstalk in CB CLI 3. Although Clastic Beanstalk still supports CB 2.6 for customers who previously installed and continue to use it, you should migrate to the latest version of CB CLI 3, as it can manage environments that you launched using CB CLI 2.6 or earlier versions of CB CLI. CB CLI automatically retrieves settings from an environment created using CB if the environment is running. Note that CB CLI 3 does not store option settings locally, as in earlier versions.

CB CLI introduces the commands eb create, eb deploy, eb open, eb console, eb scale, eb setenv, eb config, eb terminate, eb clone, eb list, eb use, eb printenv, and eb ssh. In CB CLI 3.1 or later, you can also use the eb swap command. In CB CLI 3.2 only, you can use the eb abort, eb platform, and eb upgrade commands. In addition to these new commands, CB CLI 3 commands differ from CB CLI 2.6 commands in several cases:

1. eb init - Use eb init to create an .elasticbeanstalk directory in an existing project directory and create a new Clastic Beanstalk application for the project. Unlike with previous versions, CB CLI 3 and later versions do not prompt you to create an environment.
2. eb start - CB CLI 3 does not include the command eb start. Use eb create to create an environment.
3. eb stop - CB CLI 3 does not include the command eb stop. Use eb terminate to completely terminate an environment and clean up.
4. eb push and git aws.push - CB CLI 3 does not include the commands eb push or git aws.push. Use eb deploy to update your application code.
5. eb update - CB CLI 3 does not include the command eb update. Use eb config to update an environment.
6. eb branch - CB CLI 3 does not include the command eb branch.

For more information about using CB CLI 3 commands to create and manage an application, see CB CLI Command Reference. For a command reference for CB 2.6, see CB CLI 2 Commands. For a walkthrough of how to deploy a sample application using CB CLI 3, see Managing Clastic Beanstalk environments with the CB CLI. For a walkthrough of how to deploy a sample application using eb 2.6, see Getting Started with Cb. For a walkthrough of how to use CB 2.6 to map a Git branch to a specific environment, see Deploying a Git Branch to a Specific environment. <https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/eb-cli.html#eb-cli2-differences> Note: Additionally, CB CLI 2.6 has been deprecated. It has been replaced by AWS CLI <https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/eb-cli3.html> I We will replace this question soon.

NEW QUESTION 188

Your company is supporting a number of applications that need to be moved to AWS. Initially they thought of moving these applications to the Elastic beanstalk service. When going to the Elastic beanstalk service, you can see that the underlying platform service is not an option in the Elastic beanstalk service. Which of the following options can be used to port your application onto Elastic beanstalk

- A. Use the Opswork service to create a stack
- B. In the stack, create a separate custom layer
- C. Deploy the application to this layer and then attach the layer to Elastic beanstalk
- D. Use custom chef recipe's to deploy your application in Elastic beanstalk.
- E. Use custom CloudFormation templates to deploy the application into Elastic beanstalk
- F. Create a Docker container for the custom application and then deploy it to Elastic beanstalk.

Answer: D

Explanation:

The AWS documentation mentions the following

Clastic 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.

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

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

NEW QUESTION 193

Which of the following services can be used to detect the application health in a Blue Green deployment in AWS?

- A. AWSCode Commit
- B. AWSCode Pipeline
- C. AWSCloudTrail
- D. AWSCloudwatch

Answer: D

Explanation:

The AWS Documentation mentions the following

Amazon Cloud Watch is a monitoring service for AWS Cloud resources and the applications you run on AWS. CloudWatch can collect and track metrics, collect and monitor log files, and set alarms. It provides system-wide visibility into resource utilization, application performance, and operational health, which are key to early detection of application health in blue/green deployments.

For more information on Blue Green deployments, please refer to the below link:

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

NEW QUESTION 196

Your company has an application hosted in AWS which makes use of DynamoDB. There is a requirement from the IT security department to ensure that all source IP addresses which make calls to the DynamoDB tables are recorded. Which of the following services can be used to ensure this requirement is fulfilled.

- A. AWSCode Commit
- B. AWSCode Pipeline
- C. AWSCloudTrail
- D. AWSCloudwatch

Answer: C

Explanation:

The AWS Documentation mentions the following

DynamoDB is integrated with CloudTrail, a service that captures low-level API requests made by or on behalf of DynamoDB in your AWS account and delivers the log files to an Amazon S3 bucket that you specify. CloudTrail captures calls made from the DynamoDB console or from the DynamoDB low-level API. Using the information collected by CloudTrail, you can determine what request was made to DynamoDB, the source IP address from which the request was made, who made the request, when it was made, and so on.

For more information on DynamoDB and Cloudtrail, please refer to the below link:

- <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/logging-using-cloudtrail.html>

NEW QUESTION 199

Your company has a set of EC2 Instances that access data objects stored in an S3 bucket. Your IT Security department is concerned about the security of this architecture and wants you to implement the following

- 1) Ensure that the EC2 Instance securely accesses the data objects stored in the S3 bucket
- 2) Ensure that the integrity of the objects stored in S3 is maintained.

Which of the following would help fulfil the requirements of the IT Security department. Choose 2 answers from the options given below

- A. Create an IAM user and ensure the EC2 Instances uses the IAM user credentials to access the data in the bucket.
- B. Create an IAM Role and ensure the EC2 Instances uses the IAM Role to access the data in the bucket.
- C. Use S3 Cross Region replication to replicate the objects so that the integrity of data is maintained.
- D. Use an S3 bucket policy that ensures that MFA Delete is set on the objects in the bucket

Answer: BD

Explanation:

The AWS Documentation mentions the following

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 IAM Roles, please refer to the below link:

- <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

MFA Delete can be used to add another layer of security to S3 Objects to prevent accidental deletion of objects. For more information on MFA Delete, please refer to the below link:

- <https://aws.amazon.com/blogs/security/securing-access-to-aws-using-mfa-part-3/>

NEW QUESTION 204

You have a set of EC2 Instances hosting an nginx server and a web application that is used by a set of users in your organization. After a recent application version upgrade, the instance runs into technical issues and needs an immediate restart. This does not give you enough time to inspect the cause of the issue on the server. Which of the following options if implemented prior to the incident would have assisted in detecting the underlying cause of the issue?

- A. Enable detailed monitoring and check the Cloudwatch metrics to see the cause of the issue.
- B. Create a snapshot of the EBS volume before restart, attach it to another instance as a volume and then diagnose the issue.
- C. Stream all the data to Amazon Kinesis and then analyze the data in real time.
- D. Install Cloudwatch logs agent on the instance and send all the logs to Cloudwatch logs.

Answer: D

Explanation:

The AWS documentation mentions the following

You can publish log data from Amazon EC2 instances running Linux or Windows Server, and logged events from AWS CloudTrail. CloudWatch Logs can consume logs

from resources in any region, but you can only view the log data in the CloudWatch console in the regions where CloudWatch Logs is supported.

Option A is invalid as detailed monitoring will only help us to get more information about the performance metrics of the instances, volumes etc and will not be able to provide full information regarding technical issues.

Option B is incorrect if we had created a snapshot prior to the update it might be useful but not after the incident.

Option C is incorrect here we are dealing with an issue concerning the underlying application that handles the data so this solution will not help.

For more information on Cloudwatch logs, please refer to the below link:

- <http://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/StartTheCW LAgent.htm> I

NEW QUESTION 205

Your company is hosting an application in AWS. The application consists of a set of web servers and AWS RDS. The application is a read intensive application. It has been noticed that the response time of the application decreases due to the load on the AWS RDS instance. Which of the following measures can be taken to scale the data tier. Choose 2 answers from the options given below

- A. CreateAmazon DB Read Replica'
- B. Configure the application layer to query the readreplica's for query needs.
- C. UseAutoscaling to scale out and scale in the database tier
- D. UseSQS to cache the database queries
- E. UseElastiCache in front of your Amazon RDS DB to cache common queries.

Answer: AD

Explanation:

The AWS documentation mentions the following

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.

For more information on AWS RDS Read Replica's, please visit the below URL:

? <https://aws.amazon.com/rds/details/read-replicas/>

Amazon OastiCache is a web service that makes it easy to deploy, operate, and scale an in-memory data store or cache in the cloud. The service improves the performance of web applications by allowing you to retrieve information from fast, managed, in-memory data stores, instead of relying entirely on slower disk-based databases.

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

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

NEW QUESTION 207

You have a requirement to automate the creation of EBS Snapshots. Which of the following can be used to achieve this in the best way possible?

- A. Createa powershell script which uses the AWS CLI to get the volumes and then run thescript as a cron job.
- B. Usethe A WSConf ig service to create a snapshot of the AWS Volumes
- C. Usethe AWS CodeDeploy service to create a snapshot of the AWS Volumes
- D. UseCloudwatch Events to trigger the snapshots of EBS Volumes

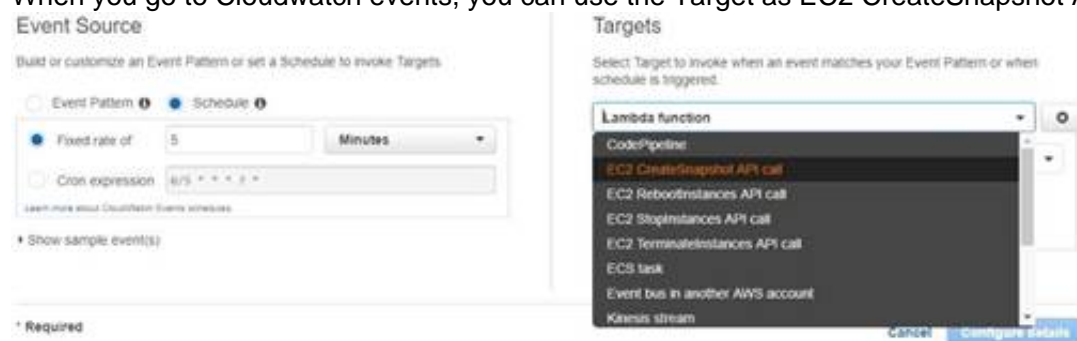
Answer: D

Explanation:

The best is to use the inbuilt sen/ice from Cloudwatch, as Cloud watch Events to automate the creation of CBS Snapshots. With Option A, you would be restricted to

running the powrshell script on Windows machines and maintaining the script itself And then you have the overhead of having a separate instance just to run that script.

When you go to Cloudwatch events, you can use the Target as EC2 CreateSnapshot API call as shown below.



The AWS Documentation mentions

Amazon Cloud Watch Cvents delivers a near real-time stream of system events that describe changes in Amazon Web Services (AWS) resources. Using simple rules that you can quickly set up, you can match events and route them to one or more target functions or streams. Cloud Watch Cvents becomes aware of operational changes as they occur. Cloud Watch Cvents responds to these operational changes and takes corrective action as necessary, by sending messages to respond to the environment, activating functions, making changes, and capturing state information.

For more information on Cloud watch Cvents, please visit the below U RL:

- <http://docs.aws.amazon.com/AmazonCloudWatch/latest/events/WhatIsCloudWatchCvents.html>

NEW QUESTION 208

When you add lifecycle hooks to an Autoscaling Group, what are the wait states that occur during the scale in and scale out process. Choose 2 answers from the options given below

- A. Launching:Wait
- B. Exiting:Wait
- C. Pending:Wait
- D. Terminating:Wait

Answer: CD

Explanation:

The AWS Documentation mentions the following

After you add lifecycle hooks to your Auto Scaling group, they work as follows:

1. Auto Scaling responds to scale out events by launching instances and scale in events by terminating instances.
2. Auto Scaling puts the instance into a wait state (Pending:Wait or Terminating: Wait). The instance is paused until either you tell Auto Scaling to continue or the timeout period ends.

For more information on Autoscaling Lifecycle hooks, please visit the below URL: • <http://docs.aws.amazon.com/autoscaling/latest/userguide/lifecycle-hooks.html>

NEW QUESTION 210

Which of the following items are required to allow an application deployed on an EC2 instance to write data to a DynamoDB table? Assume that no security keys are allowed to be stored on the EC2 instance. Choose 2 answers from the options below

- A. Create an IAM Role that allows write access to the DynamoDB table.
- B. Add an IAM Role to a running EC2 instance.
- C. Create an IAM User that allows write access to the DynamoDB table.
- D. Add an IAM User to a running EC2 instance.

Answer: AB

Explanation:

The AWS documentation mentions the following

We designed IAM roles so that your applications can securely make API requests from your instances, without requiring you to manage the security credentials that

the applications use. Instead of creating and distributing your AWS credentials, you can delegate permission to make API requests using IAM roles

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

NEW QUESTION 214

You have a set of EC2 Instances running behind an ELB. These EC2 Instances are launched via an Autoscaling Group. There is a requirement to ensure that the logs from the server are stored in a durable storage layer. This is so that log data can be analyzed by staff in the future. Which of the following steps can be implemented to ensure this requirement is fulfilled. Choose 2 answers from the options given below

- A. On the web servers, create a scheduled task that executes a script to rotate and transmit the logs to an Amazon S3 bucket
- B. *
- C. Use AWS Data Pipeline to move log data from the Amazon S3 bucket to Amazon Redshift in order to process and run reports
- D. On the web servers, create a scheduled task that executes a script to rotate and transmit the logs to Amazon Glacier.
- E. Use AWS Data Pipeline to move log data from the Amazon S3 bucket to Amazon SQS in order to process and run reports

Answer: AB

Explanation:

Amazon S3 is the perfect option for durable storage. The AWS Documentation mentions the following on S3 Storage Amazon Simple Storage Service (Amazon S3) makes it simple and practical to collect, store, and analyze data - regardless of format - all at massive scale. S3 is object storage built to store and retrieve any amount of data from anywhere - web sites and mobile apps, corporate applications, and data from IoT sensors or devices.

For more information on Amazon S3, please refer to the below URL:

- <https://aws.amazon.com/s3/>

Amazon Redshift is a fast, fully managed data warehouse that makes it simple and cost-effective to analyze all your data using standard SQL and your existing Business Intelligence (BI) tools. It allows you to run complex analytic queries against petabytes of structured data, using sophisticated query optimization, columnar storage on high-performance local disks, and massively parallel query execution. Most results come back in seconds. For more information on Amazon Redshift, please refer to the below URL:

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

NEW QUESTION 217

Your company has a set of development teams that work in a variety of programming languages. They develop applications which have a lot of different application dependencies. There is a move from the company to move these development environments onto AWS. Which of the below is the best option to make this happen.

- A. Use the CloudFormation service to create Docker containers for each type of application
- B. Use the Elastic Beanstalk service and use Docker containers to host each application environment for the developer community
- C. Use the OpsWorks service, create a stack and create separate layers for each application environment for the developer community
- D. Launch separate EC2 Instances to host each application type for the developer community

Answer: B

Explanation:

The AWS Documentation mentions the following

Elastic Beanstalk supports the deployment of web applications from Docker containers. With Docker containers, you can define your own runtime environment. You can choose your own platform, programming language, and any application dependencies (such as package managers or tools), that aren't supported by other platforms. Docker containers are self-contained and include all the configuration information and software your web application requires to run.

For more information on Elastic Beanstalk and Docker, please refer to the below URL:

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

NEW QUESTION 219

You are setting up CloudFormation templates for your organization. The CloudFormation template consists of creating EC2 Instances for both your development and production environments in the same region. Each of these instances will have an Elastic IP and a security group attached to them which will be done via CloudFormation. Your CloudFormation stack for the development environment gets successfully created, but then the production CloudFormation stack fails. Which of the below could be a reason for this.

- A. You have chosen the wrong tags when creating the instances in both environments.
- B. You hit the soft limit of 5 EIPs per region when creating the development environment.
- C. You hit the soft limit for security groups when creating the development environment.

D. You didn't choose the Production version of the AMI you are using when creating the production stack.

Answer: B

Explanation:

The most viable reason could be that you reached the limit for the number of Elastic IP's in the region.

For more information on AWS EC2 service limits, please refer to the below link:

- <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-resource-limits.html>

NEW QUESTION 222

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

DOP-C01 Practice Exam Features:

- * DOP-C01 Questions and Answers Updated Frequently
- * DOP-C01 Practice Questions Verified by Expert Senior Certified Staff
- * DOP-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * DOP-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The DOP-C01 Practice Test Here](#)