Amazon

Exam Questions AWS-Certified-DevOps-Engineer-Professional

Amazon AWS Certified DevOps Engineer Professional

NEW QUESTION 1
You are building a game high score table in DynamoDB. You will store each user's highest score for each game, with many games, all of which have relatively similar usage levels and numbers of players. You need to be able to look up the highest score for any game. What's the best DynamoDB key structure?

A. HighestScore as the hash / only key.
B. GameID as the hash key, HighestScore as the range key.
C. GameID as the hash / only key.
D. GameID as the range / only ke

Answer: B

Explanation:
Since access and storage for games is uniform, and you need to have ordering within each game for the scores (to access the highest value), your hash (partition) key should be the GameID, and there should be a range key for HighestScore.
Reference: http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GuidelinesForTables.html#Guideli nesForTables.Partitions

NEW QUESTION 2
You need to scale an RDS deployment. You are operating at 10% writes and 90% reads, based on your logging. How best can you scale this in a simple way?

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

Answer: C

Explanation:
The high-availability feature is not a scaling solution for read-only scenarios; you cannot use a standby replica to serve read traffic. To service read-only traffic, you should use a Read Replica. For more information, see Working with PostgreSQL, MySQL, and NIariaDB Read Replicas.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.NIuItiAZ.htmI

NEW QUESTION 3
When thinking of AWS Elastic Beanstalk, the 'Swap Environment URLs' feature most directly aids in what?

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

Answer: D

Explanation:
Simply upload the new version of your application and let your deployment service (AWS Elastic Beanstalk, AWS CloudFormation, or AWS OpsWorks) deploy a new version (green). To cut over to the new version, you simply replace the ELB URLs in your DNS records. Elastic Beanstalk has a Swap Environment URLs feature to facilitate a simpler cutover process.
Reference: https://d0.awsstatic.com/whitepapers/overview-of-deployment-options-on-aws.pdf

NEW QUESTION 4
You need to process long-running jobs once and only once. How might you do this?

A. Use an SNS queue and set the visibility timeout to long enough forjobs to process.
B. Use an SQS queue and set the reprocessing timeout to long enough forjobs to process.
C. Use an SQS queue and set the visibility timeout to long enough forjobs to process.
D. Use an SNS queue and set the reprocessing timeout to long enough forjobs to proces

Answer: C

Explanation:
The message timeout defines how long after a successful receive request SQS waits before allowing jobs to be seen by other components, and proper configuration prevents duplicate processing.
Reference: http://docs.aws.amazon.com/AWSSimpIeQueueService/latest/SQSDeveIoperGuide/MessageLifecycle.ht ml

NEW QUESTION 5
Which of these is not an instrinsic function in AWS CloudFormation?

A. Fn::Equals
B. Fn::|f
C. Fn::Not
D. Fn::Parse

Answer: D

Explanation:
This is the complete list of Intrinsic Functions...: Fn::Base64, Fn::And, Fn::Equals, Fn::If, Fn::Not, Fn::Or, Fn::FindInMap, Fn::GetAtt, Fn::GetAZs, Fn::Join, Fn::Select, Ref
Reference:

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

NEW QUESTION 6
For AWS Auto Scaling, what is the first transition state an instance enters after leaving steady state when scaling in due to health check failure or decreased load?

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

Answer: A

Explanation:
When Auto Scaling responds to a scale in event, it terminates one or more instances. These instances are detached from the Auto Scaling group and enter the Terminating state.
Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AutoScalingGroupLifecycle.html

NEW QUESTION 7
You need to create a simple, holistic check for your system's general availablity and uptime. Your system presents itself as an HTTP-speaking API. What is the most simple tool on AWS to achieve this with?

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

Answer: A

Explanation:
You can create a health check that will run into perpetuity using Route53, in one API call, which will ping your service via HTTP every 10 or 30 seconds.
Amazon Route 53 must be able to establish a TCP connection with the endpoint within four seconds. In addition, the endpoint must respond with an HTTP status code of 200 or greater and less than 400 within two seconds after connecting.
Reference:
http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover-determining-health-of-endpoint s.html

NEW QUESTION 8
What is the scope of an EC2 EIP?

A. Placement Group
B. Availability Zone
C. Region
D. VPC

Answer: C

Explanation:
An Elastic IP address is tied to a region and can be associated only with an instance in the same region. Reference:
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/resources.html

NEW QUESTION 9
You want to pass queue messages that are 1GB each. How should you achieve this?

A. Use Kinesis as a buffer stream for message bodie
B. Store the checkpoint id for the placement in the Kinesis Stream in SQS.
C. Use the Amazon SQS Extended Client Library for Java and Amazon S3 as a storage mechanism for message bodies.
D. Use SQS's support for message partitioning and multi-part uploads on Amazon S3.
E. Use AWS EFS as a shared pool storage mediu
F. Store filesystem pointers to the files on disk in the SQS message bodies.

Answer: B

Explanation:
You can manage Amazon SQS messages with Amazon S3. This is especially useful for storing and retrieving messages with a message size of up to 2 GB. To manage Amazon SQS messages with Amazon S3, use the Amazon SQS Extended Client Library for Java.
Reference:
http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/s3-messages.html

NEW QUESTION 10
There are a number of ways to purchase compute capacity on AWS. Which orders the price per compute or memory unit from LOW to HIGH (cheapest to most expensive), on average?
(A) On-Demand (B) Spot (C) Reserved

A. A, B, C
B. C, B, A
C. B, C, A
D. A, C, B

Answer:

C

Explanation:
Spot instances are usually many, many times cheaper than on-demand prices. Reserved instances, depending on their term and utilization, can yield approximately 33% to 66% cost savings. On-Demand prices are the baseline price and are the most expensive way to purchase EC2 compute time. Reference: https://d0.awsstatic.com/whitepapers/Cost_Optimization_with_AWS.pdf

NEW QUESTION 10
When thinking of DynamoDB, what are true of Local Secondary Key properties?

A. Either the partition key or the sort key can be different from the table, but not both.
B. Only the sort key can be different from the table.
C. The partition key and sort key can be different from the table.
D. Only the partition key can be different from the tabl

Answer: B

Explanation:
Global secondary index — an index with a partition key and a sort key that can be different from those on the table. A global secondary index is considered "global" because queries on the index can span all of the data in a table, across all partitions.
Reference: http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/SecondaryIndexes.html

NEW QUESTION 13
Which deployment method, when using AWS Auto Scaling Groups and Auto Scaling Launch Configurations, enables the shortest time to live for indMdual sewers?

A. Pre-baking AMIs with all code and configuration on deploys.
B. Using a Dockerfile bootstrap on instance launch.
C. Using UserData bootstrapping scripts.
D. Using AWS EC2 Run Commands to dynamically SSH into fileet

Answer: A

Explanation:
Note that the bootstrapping process can be slower if you have a complex application or multiple applications to install. Managing a fileet of applications with several build tools and dependencies can be a challenging task during rollouts. Furthermore, your deployment service should be designed to do faster rollouts to take advantage of Auto Scaling. Prebaking is a process of embedding a significant portion of your application artifacts within your base AMI. During the deployment process you can customize application installations by using EC2 instance artifacts such as instance tags, instance metadata, and Auto Scaling groups.
Reference: https://d0.awsstatic.com/whitepapers/overview-of-deployment-options-on-aws.pdf

NEW QUESTION 18
Which of these techniques enables the fastest possible rollback times in the event of a failed deployment?

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

Answer: D

Explanation:
AWS specifically recommends Blue-Green for super-fast, zero-downtime deploys - and thus rollbacks, which are redeploying old code.
You use various strategies to migrate the traffic from your current application stack (blue) to a new version of the application (green). This is a popular technique for deploying applications with zero downtime. Reference: https://d0.awsstatic.com/whitepapers/overview-of-deployment-options-on-aws.pdf

NEW QUESTION 23
Which major database needs a BYO license?

A. PostgreSQL
B. NlariaDB
C. MySQL
D. Oracle

Answer: D

Explanation:
Oracle is not open source, and requires a bring your own license model.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Oracle.html|

NEW QUESTION 26
For AWS Auto Scaling, what is the first transition state a new instance enters after leaving steady state when scaling out due to increased load?

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

Answer: B

Explanation:
When a scale out event occurs, the Auto Scaling group launches the required number of EC2 instances, using its assigned launch configuration. These instances start in the Pending state. If you add a lifecycle hook to your Auto Scaling group, you can perform a custom action here. For more information, see Lifecycle Hooks.
Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AutoScalingGroupLifecycle.html

NEW QUESTION 31
You need to know when you spend $1000 or more on AWS. What's the easy way for you to see that notification?

A. AWS CIoudWatch Events tied to API calls, when certain thresholds are exceeded, publish to SNS.
B. Scrape the billing page periodically and pump into Kinesis.
C. AWS CIoudWatch Metrics + Billing Alarm + Lambda event subscriptio
D. When a threshold is exceeded, email the manager.
E. Scrape the billing page periodically and publish to SN

Answer: C

Explanation:
Even if you're careful to stay within the free tier, it's a good idea to create a billing alarm to notify you if you exceed the limits of the free tier. Billing alarms can help to protect you against unknowingly accruing charges if you inadvertently use a service outside of the free tier or if traffic exceeds your expectations. Reference: http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/free-tier-alarms.html

NEW QUESTION 33
Which of these is not a Pseudo Parameter in AWS CIoudFormation?

A. AWS::StackName
B. AWS::AccountId
C. AWS::StackArn
D. AWS::NotificationARNs

Answer: C

Explanation:
This is the complete list of Pseudo Parameters: AWS::Account|d, AWS::NotificationARNs, AWS::NoValue, AWS::Region, AWS::StackId, AWS::StackName
Reference:
http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/pseudo-parameter-reference.html

NEW QUESTION 38
What is the scope of an EBS volume?

A. VPC
B. Region
C. Placement Group
D. Availability Zone

Answer: D

Explanation:
An Amazon EBS volume is tied to its Availability Zone and can be attached only to instances in the same Availability Zone.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/resources.html

NEW QUESTION 40
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:
standard volumes, or Magnetic volumes, are best for: Cold workloads where data is infrequently accessed, or scenarios where the lowest storage cost is important.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html

NEW QUESTION 44
You need to deploy an AWS stack in a repeatable manner across multiple environments. You have selected CIoudFormation as the right tool to accomplish this, but have found that there is a resource type you need to create and model, but is unsupported by CIoudFormation. How should you overcome this chaHenge?

A. Use a CIoudFormation Custom Resource Template by selecting an API call to proxy for create, update, and delete action

B. CloudFormation will use the AWS SDK, CLI, or API method of your choosing as the state transition function for the resource type you are modeling.
C. Submit a ticket to the AWS Forum
D. AWS extends CloudFormation Resource Types by releasing tooling to the AWS Labs organization on GitHu
E. Their response time is usually 1 day, and they complete requests within a week or two.
F. Instead of depending on CloudFormation, use Chef, Puppet, or Ansible to author Heat templates, which are declarative stack resource definitions that operate over the OpenStack hypervisor and cloud environment.
G. Create a CloudFormation Custom Resource Type by implementing create, update, and delete functionality, either by subscribing a Custom Resource Provider to an SNS topic, or by implementing the logic in AWS Lambda.

Answer: D

Explanation:
Custom resources provide a way for you to write custom provisioning logic in AWS CloudFormation template and have AWS CloudFormation run it during a stack operation, such as when you create, update or delete a stack. For more information, see Custom Resources.
Reference:
http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/template-custom-resources.html

NEW QUESTION 45
You run a 2000-engineer organization. You are about to begin using AWS at a large scale for the first time. You want to integrate with your existing identity management system running on Microsoft Active Directory, because your organization is a power-user of Active Directory. How should you manage your AWS identities in the most simple manner?

A. Use a large AWS Directory Service Simple AD.
B. Use a large AWS Directory Service AD Connector.
C. Use an Sync Domain running on AWS Directory Service.
D. Use an AWS Directory Sync Domain running on AWS Lambda

Answer: B

Explanation:
You must use AD Connector as a power-user of Microsoft Active Directory. Simple AD only works with a subset of AD functionality. Sync Domains do not exist; they are made up answers.
AD Connector is a directory gateway that allows you to proxy directory requests to your on-premises Nlicrosoft Active Directory, without caching any information in the cloud. AD Connector comes in 2 sizes; small and large. A small AD Connector is designed for smaller organizations of up to 500 users. A large AD Connector is designed for larger organizations of up to 5,000 users.
Reference: https://aws.amazon.com/directoryservice/details/

NEW QUESTION 48
When thinking of AWS OpsWorks, which of the following is not an instance type you can allocate in a stack layer?

A. 24/7 instances
B. Spot instances
C. Time-based instances
D. Load-based instances

Answer: B

Explanation:
AWS OpsWorks supports the following instance types, which are characterized by how they are started and stopped. 24/7 instances are started manually and run until you stop them.Time-based instances are run by AWS OpsWorks on a specified daily and weekly schedule. They allow your stack to automatically adjust the number of instances to accommodate predictable usage patterns. Load-based instances are automatically started and stopped by AWS OpsWorks, based on specified load metrics, such as CPU utilization. They allow your stack to automatically adjust the number of instances to accommodate variations in incoming traffic. Load-based instances are available only for Linux-based stacks. Reference: http://docs.aws.amazon.com/opsworks/latest/userguide/welcome.html

NEW QUESTION 53
You are building a Ruby on Rails application for internal, non-production use which uses IV|ySQL as a database. You want developers without very much AWS experience to be able to deploy new code with a single command line push. You also want to set this up as simply as possible. Which tool is ideal for this setup?

A. AWS CloudFormation
B. AWS OpsWorks
C. AWS ELB + EC2 with CLI Push
D. AWS Elastic Beanstalk

Answer: D

Explanation:
Elastic Beanstalk's primary mode of operation exactly supports this use case out of the box. It is simpler than all the other options for this question.
With Elastic Beanstalk, you can quickly deploy and manage applications in the AWS cloud without worrying about the infrastructure that runs those applications.
AWS Elastic Beanstalk reduces management complexity without restricting choice or control. You simply upload your application, and Elastic Beanstalk automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring.
Reference: http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_Ruby_rails.html

NEW QUESTION 57
You are building a mobile app for consumers to post cat pictures online. You will be storing the images in AWS S3. You want to run the system very cheaply and simply. Which one of these options allows you to build a photo sharing application without needing to worry about scaling expensive uploads processes, authentication/authorization and so forth?

A. Build the application out using AWS Cognito and web identity federation to allow users to log in using Facebook or Google Account
B. Once they are logged in, the secret token passed to that user is used to directly access resources on AWS, like AWS S3.

C. Use JWT or SANIL compliant systems to build authorization policie
D. Users log in with a username and password, and are given a token they can use indefinitely to make calls against the photo infrastructure.
E. Use AWS API Gateway with a constantly rotating API Key to allow access from the client-sid
F. Construct a custom build of the SDK and include S3 access in it.
G. Create an AWS oAuth Service Domain ad grant public signup and access to the domai
H. During setup, add at least one major social media site as a trusted Identity Provider for users.

Answer: A

Explanation:
The short answer is that Amazon Cognito is a superset of the functionality provided by web identity federation. It supports the same providers, and you configure your app and authenticate with those providers in the same way. But Amazon Cognito includes a variety of additional features. For example, it enables your users to start using the app as a guest user and later sign in using one of the supported identity providers.
Reference:
https://blogs.aws.amazon.com/security/post/Tx3SYCORF5EKRCO/How-Does-Amazon-Cognito-Relate-to
-Existing-Web-Identity-Federatio

NEW QUESTION 62
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 Trai
B. Configure a script to aggregate the log data delivered to S3 once per week and deliver this to the CTO.
C. Use CloudWatch Events Rules with an SNS topic subscribed to all AWS API call
D. Subscribe the CTO to an email type delivery on this SNS Topic.
E. Use AWS IAM credential reports to deliver a CSV of all uses of IAM User Tokens over time to the CTO.
F. Use AWS Config with an SNS subscription on a Lambda, and insert these changes over time into a DynamoDB tabl
G. Generate reports based on the contents of this table.

Answer: A

Explanation:
This is the ideal use case for AWS CloudTrail.
CloudTrail provides visibility into user actMty by recording API calls made on your account. CloudTrail records important information about each API call, including the name of the API, the identity of the caller, the time of the API call, the request parameters, and the response elements returned by the AWS service. This information helps you to track changes made to your AWS resources and to troubleshoot operational issues. CloudTrail makes it easier to ensure compliance with internal policies and regulatory standards. Reference: https://aws.amazon.com/Cloudtrail/faqs/

NEW QUESTION 65
Which is not a restriction on AWS EBS Snapshots?

A. Snapshots which are shared cannot be used as a basis for other snapshots.
B. You cannot share a snapshot containing an AWS Access Key ID or AWS Secret Access Key.
C. You cannot share unencrypted snapshots.
D. Snapshot restorations are restricted to the region in which the snapshots are create

Answer: A

Explanation:
Snapshots shared with other users are usable in full by the recipient, including but limited to the ability to base modified volumes and snapshots.
Reference:
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-modifying-snapshot-permissions.html

NEW QUESTION 69
If I want CloudFormation stack status updates to show up in a continuous delivery system in as close to real time as possible, how should I achieve this?

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

Answer: C

Explanation:
Use NotificationARNs.member.N when making a CreateStack call to push stack events into SNS in nearly real-time.
Reference:
http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-updating-stacks-monitor-s tack.html

NEW QUESTION 71
What does it mean if you have zero IOPS and a non-empty I/O queue for all EBS volumes attached to a running EC2 instance?

A. The I/O queue is buffer flushing.
B. Your EBS disk head(s) is/are seeking magnetic stripes.
C. The EBS volume is unavailable.
D. You need to re-mount the EBS volume in the O

Answer: C

Explanation:

This is the definition of Unavailable from the EC2 and EBS SLA.
"Unavailable" and "Unavai|abi|ity" mean... For Amazon EBS, when all of your attached volumes perform zero read write IO, with pending IO in the queue.
Reference: https://aws.amazon.com/ec2/s|a/

NEW QUESTION 75
Which of these configuration or deployment practices is a security risk for RDS?

A. Storing SQL function code in plaintext
B. Non-Multi-AZ RDS instance
C. Having RDS and EC2 instances exist in the same subnet
D. RDS in a public subnet

Answer: D

Explanation:
Making RDS accessible to the public internet in a public subnet poses a security risk, by making your database directly addressable and spammable.
DB instances deployed within a VPC can be configured to be accessible from the Internet or from EC2 instances outside the VPC. If a VPC security group specifies a port access such as TCP port 22, you would not be able to access the DB instance because the firewall for the DB instance provides access only via the IP addresses specified by the DB security groups the instance is a member of and the port defined when the DB instance was created.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.RDSSecurityGroups.htmI

NEW QUESTION 79
Which of these is not a reason a Mu|ti-AZ RDS instance will failover?

A. An Availability Zone outage
B. A manual failover of the DB instance was initiated using Reboot with failover
C. To autoscale to a higher instance class
D. The primary DB instance fails

Answer: C

Explanation:
The primary DB instance switches over automatically to the standby replica if any of the > following conditions occur: An Availability Zone outage, the primary DB instance fails, the DB instance's server type is changed, the operating system of the DB instance is, undergoing software patching, a manual failover of the DB instance was initiated using Reboot with failover
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.Mu|tiAZ.htmI

NEW QUESTION 82
You need your API backed by DynamoDB to stay online during a total regional AWS failure. You can tolerate a couple minutes of lag or slowness during a large failure event, but the system should recover with normal operation after those few minutes. What is a good approach?

A. Set up DynamoDB cross-region replication in a master-standby configuration, with a single standby in another regio
B. Create an Auto Scaling Group behind an ELB in each of the two regions DynamoDB is running i
C. Add a Route53 Latency DNS Record with DNS Failover, using the ELBs in the two regions as the resource records.
D. Set up a DynamoDB Multi-Region tabl
E. Create an Auto Scaling Group behind an ELB in each of the two regions DynamoDB is running i
F. Add a Route53 Latency DNS Record with DNS Failover, using the ELBs in the two regions as the resource records.
G. Set up a DynamoDB Mu|ti-Region tabl
H. Create a cross-region ELB pointing to a cross-region Auto Scaling Group, and direct a Route53 Latency DNS Record with DNS Failover to the cross-region ELB.
I. Set up DynamoDB cross-region replication in a master-standby configuration, with a single standby in another regio
J. Create a cross-region ELB pointing to a cross-region Auto Scaling Group, and direct a Route53 Latency DNS Record with DNS Failover to the cross-region ELB.

Answer: A

Explanation:
There is no such thing as a cross-region ELB, nor such thing as a cross-region Auto Scaling Group, nor such thing as a DynamoDB Multi-Region Table. The only option that makes sense is the cross-regional replication version with two ELBs and ASGs with Route53 Failover and Latency DNS.
Reference: http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.CrossRegionRepI.htmI

NEW QUESTION 83
Your company wants to understand where cost is coming from in the company's production AWS account. There are a number of applications and services running at any given time. Without expending too much initial development time, how best can you give the business a good understanding of which applications cost the most per month to operate?

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

Answer: C

Explanation:
Cost Allocation Tagging is a built-in feature of AWS, and when coupled with the Cost Explorer, provides a simple and robust way to track expenses.
You can also use tags to filter views in Cost Explorer. Note that before you can filter views by tags in Cost Explorer, you must have applied tags to your resources and activate them, as described in the following sections. For more information about Cost Explorer, see Analyzing Your Costs with Cost Explorer. Reference:
http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/cost-alloc-tags.html

NEW QUESTION 86
Which of the following tools does not directly support AWS OpsWorks, for monitoring your stacks?

A. AWS Config
B. Amazon CloudWatch Nletrics
C. AWS CloudTrail
D. Amazon CloudWatch Logs

Answer: A

Explanation:
You can monitor your stacks in the following ways: AWS OpsWorks uses Amazon CloudWatch to provide thirteen custom metrics with detailed monitoring for each instance in the stack; AWS OpsWorks integrates with AWS CloudTrail to log every AWS OpsWorks API call and store the data in an Amazon S3 bucket; You can use Amazon CloudWatch Logs to monitor your stack's system, application, and custom logs. Reference:
http://docs.aws.amazon.com/opsworks/latest/userguide/monitoring.html

NEW QUESTION 90
When thinking of AWS OpsWorks, which of the following is true?

A. Stacks have many layers, layers have many instances.
B. Instances have many stacks, stacks have many layers.
C. Layers have many stacks, stacks have many instances.
D. Layers have many instances, instances have many stack

Answer: A

Explanation:
The stack is the core AWS OpsWorks component. It is basically a container for AWS resources—Amazon EC2 instances, Amazon RDS database instances, and so on—that have a common purpose and should
be logically managed together. You define the stack's constituents by adding one or more layers. A layer represents a set of Amazon EC2 instances that serve a particular purpose, such as serving applications or hosting a database server. An instance represents a single computing resource, such as an Amazon EC2 instance.
Reference: http://docs.aws.amazon.com/opsworks/latest/userguide/welcome.html

NEW QUESTION 93
You work for a company that automatically tags photographs using artificial neural networks (ANNs), which run on GPUs using C++. You receive millions of images at a time, but only 3 times per day on average. These images are loaded into an AWS S3 bucket you control for you in a batch, and then the customer publishes a JSON-formatted manifest into another S3 bucket you control as well. Each image takes 10 milliseconds to process using a full GPU. Your neural network software requires 5 minutes to bootstrap. Image tags are JSON objects, and you must publish them to an S3 bucket.
Which of these is the best system architectures for this system?

A. Create an OpsWorks Stack with two Layer
B. The first contains lifecycle scripts for launching and bootstrapping an HTTP API on G2 instances for ANN image processing, and the second has an always-on instance which monitors the S3 manifest bucket for new file
C. When a new file is detected, request instances to boot on the ANN laye
D. When the instances are booted and the HTTP APIs are up, submit processing requests to indMdual instances.
E. Make an S3 notification configuration which publishes to AWS Lambda on the manifest bucke
F. Make the Lambda create a CloudFormation Stack which contains the logic to construct an autoscaling worker tier of EC2 G2 instances with the ANN code on each instanc
G. Create an SQS queue of the images in the manifes
H. Tear the stack down when the queue is empty.
I. Deploy your ANN code to AWS Lambda as a bundled binary for the C++ extensio
J. Make an S3 notification configuration on the manifest, which publishes to another AWS Lambda running controller cod
K. This controller code publishes all the images in the manifest to AWS Kinesi
L. Your ANN code Lambda Function uses the Kinesis as an Event Sourc
M. The system automatically scales when the stream contains image events.
N. Create an Auto Scaling, Load Balanced Elastic Beanstalk worker tier Application and Environmen
O. Deploy the ANN code to G2 instances in this tie
P. Set the desired capacity to 1. Make the code periodically check S3 for new manifest
Q. When a new manifest is detected, push all of the images in the manifest into the SQS queue associated with the Elastic Beanstalk worker tier.

Answer: B

Explanation:
The Elastic Beanstalk option is incorrect because it requires a constantly-polling instance, which may break and costs money.
The Lambda fileet option is incorrect because AWS Lambda does not support GPU usage.
The OpsWorks stack option both requires a constantly-polling instance, and also requires complex timing and capacity planning logic.
The CloudFormation option requires no polling, has no always-on instances, and allows arbitrarily fast processing by simply setting the instance count as high as needed.
Reference: http://docs.aws.amazon.com/lambda/latest/dg/current-supported-versions.html

NEW QUESTION 96
You need to create a Route53 record automatically in CloudFormation when not running in production during all launches of a Template. How should you implement this?

A. Use a <code>Parameter</code> for <code>environment</code>, and add a <code>Condition</code> on the Route53 <code>Resource</code> in the template to create the record only when<code>environment</code> is not <code>production</code>.
B. Create two templates, one with the Route53 record value and one with a null value for the recor
C. Use the one without it when deploying to production.

D. Use a <code>Parameter</code> for <code>environment</code>, and add a <code>Condition</code> on the Route53 <code>Resource</code> in the template to create the record with a null string when<code>environment</code> is <code>production</code>.
E. Create two templates, one with the Route53 record and one without i
F. Use the one without it when deploying to production.

Answer: A

Explanation:
The best way to do this is with one template, and a Condition on the resource. Route53 does not allow null strings for records.
Reference:
http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/conditions-section-structure.html

NEW QUESTION 100
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 fortools and AWS resources.
Their stack: 2-tier API
Data stored in DynamoDB or S3, depending on type 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 sewers.
Which of the following approaches, given this company's stack and their priorities, best meets the company's needs?

A. Model the stack in AWS Elastic Beanstalk as a single Application with multiple Environment
B. Use Elastic Beanstalk's Rolling Deploy option to progressively roll out application code changes when promoting across environments.
C. Model the stack in 3 CloudFormation templates: Data layer, compute layer, and networking laye
D. Write stack deployment and integration testing automation following Blue-Green methodologies.
E. Model the stack in AWS OpsWorks as a single Stack, with 1 compute layer and its associated EL
F. Use Chef and App Deployments to automate Rolling Deployment.
G. Model the stack in 1 CloudFormation template, to ensure consistency and dependency graph resolutio
H. Write deployment and integration testing automation following Rolling Deployment methodologies.

Answer: B

Explanation:
AWS recommends Blue-Green for zero-downtime deploys. Since you use DynamoDB, and neither AWS OpsWorks nor AWS Elastic Beanstalk directly supports DynamoDB, the option selecting CloudFormation and Blue-Green is correct.
You use various strategies to migrate the traffic from your current application stack (blue) to a new version of the application (green). This is a popular technique for deploying applications with zero downtime. The deployment services like AWS Elastic Beanstalk, AWS CIoudFormation, or AWS OpsWorks are particularly useful as they provide a simple way to clone your running application stack. You can set up a
new version of your application (green) by simply cloning current version of the application (blue). Reference: https://d0.awsstatic.com/whitepapers/overview-of-deployment-options-on-aws.pdf

NEW QUESTION 101
When thinking of AWS Elastic Beanstalk, which statement is true?

A. Worker tiers pull jobs from SNS.
B. Worker tiers pull jobs from HTTP.
C. Worker tiers pull jobs from JSON.
D. Worker tiers pull jobs from SQ

Answer: D

Explanation:
Elastic Beanstalk installs a daemon on each Amazon EC2 instance in the Auto Scaling group to process Amazon SQS messages in the worker environment. The daemon pulls data off the Amazon SQS queue, inserts it into the message body of an HTTP POST request, and sends it to a user-configurable URL path on the local host. The content type for the message body within an HTTP POST request is application/json by default.
Reference:
http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-tiers.htmI

NEW QUESTION 102
Your application's Auto Scaling Group scales up too quickly, too much, and stays scaled when traffic decreases. What should you do to fix this?

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

Answer: B

Explanation:
Systems will always over-scale unless you choose the metric that runs out first and becomes constrained first. You also need to set the thresholds of the metric based on whether or not latency is affected by the change, tojustify adding capacity instead of wasting money.
Reference: http://docs.aws.amazon.com/AutoSca|ing/latest/DeveIoperGuide/policy_creating.htmI

NEW QUESTION 103
Your CTO is very worried about the security of your AWS account. How best can you prevent hackers from completely hijacking your account?

A. Use short but complex password on the root account and any administrators.
B. Use AWS IAM Geo-Lock and disallow anyone from logging in except for in your city.
C. Use MFA on all users and accounts, especially on the root account.
D. Don't write down or remember the root account password after creating the AWS accoun

Answer: C

Explanation:
For increased security, we recommend that you configure multi-factor authentication (MFA) to help protect your AWS resources. MFA adds extra security because it requires users to enter a unique authentication code from an approved authentication device or SMS text message when they access AWS websites or services. Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentiaIs_mfa.htmI

NEW QUESTION 107
If you're trying to configure an AWS Elastic Beanstalk worker tier for easy debugging if there are problems finishing queue jobs, what should you configure?

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

Answer: D

Explanation:
Elastic Beanstalk worker environments support Amazon Simple Queue Service (SQS) dead letter queues. A dead letter queue is a queue where other (source) queues can send messages that for some reason could not be successfully processed. A primary benefit of using a dead letter queue is the ability to sideline and isolate the unsuccessfully processed messages. You can then analyze any messages sent
to the dead letter queue to try to determine why they were not successfully processed. Reference:
http://docs.aws.amazon.com/elasticbeanstaIk/latest/dg/using-features-managing-env-tiers.htmI#worker-d eadletter

NEW QUESTION 111
......

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

AWS-Certified-DevOps-Engineer-Professional Practice Exam Features:

* AWS-Certified-DevOps-Engineer-Professional Questions and Answers Updated Frequently

* AWS-Certified-DevOps-Engineer-Professional Practice Questions Verified by Expert Senior Certified Staff

* AWS-Certified-DevOps-Engineer-Professional Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* AWS-Certified-DevOps-Engineer-Professional Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
Order The AWS-Certified-DevOps-Engineer-Professional Practice Test Here