Amazon

Exam Questions AWS-Certified-DevOps-Engineer-Professional

Amazon AWS Certified DevOps Engineer Professional

NEW QUESTION 1
You run a clustered NoSQL database on AWS EC2 using AWS EBS. You need to reduce latency for database response times. Performance is the most important concern, not availability. You did not perform the initial setup, someone without much AWS knowledge did, so you are not sure if they configured everything optimally. Which of the following is NOT likely to be an issue contributing to increased latency?

A. The EC2 instances are not EBS Optimized.
B. The database and requesting system are both in the wrong Availability Zone.
C. The EBS Volumes are not using PIOPS.
D. The database is not running in a placement grou

Answer: B

Explanation:
For the highest possible performance, all instances in a clustered database like this one should be in a single Availability Zone in a placement group, using EBS optimized instances, and using PIOPS SSD EBS Volumes. The particular Availability Zone the system is running in should not be important, as long as it is the same as the requesting resources.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html

NEW QUESTION 2
Fill the blanks: helps us track AWS API calls and transitions, helps to understand what resources we have now, and allows auditing credentials and logins.

A. AWS Config, CIoudTraiI, IAM Credential Reports
B. CIoudTraiI, IAM Credential Reports, AWS Config
C. CIoudTraiI, AWS Config, IAM Credential Reports
D. AWS Config, IAM Credential Reports, CIoudTraiI

Answer: C

Explanation:
You can use AWS CIoudTraiI to get a history of AWS API calls and related events for your account. This includes calls made by using the AWS Management Console, AWS SDKs, command line tools, and higher-level AWS services.
Reference: http://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-user-guide.html

NEW QUESTION 3
Your system automatically provisions EIPs to EC2 instances in a VPC on boot. The system provisions the whole VPC and stack at once. You have two of them per VPC. On your new AWS account, your attempt to create a Development environment failed, after successfully creating Staging and Production environments in the same region. What happened?

A. You didn't choose the Development version of the AMI you are using.
B. You didn't set the Development flag to true when deploying EC2 instances.
C. You hit the soft limit of 5 EIPs per region and requested a 6th.
D. You hit the soft limit of 2 VPCs per region and requested a 3r

Answer: C

Explanation:
There is a soft limit of 5 E|Ps per Region for VPC on new accounts. The third environment could not allocate the 6th EIP.
Reference: http://docs.aws.amazon.com/generaI/latest/gr/aws_service_|imits.htmI#Iimits_vpc

NEW QUESTION 4
To monitor API calls against our AWS account by different users and entities, we can use to create a history of calls in bulk for later review, and use for reacting to AWS API calls in real-time.

A. AWS Config; AWS Inspector
B. AWS CIoudTraiI; AWS Config
C. AWS CIoudTraiI; CIoudWatch Events
D. AWS Config; AWS Lambda

Answer: C

Explanation:
CIoudTraiI is a batch API call collection service, CIoudWatch Events enables real-time monitoring of calls through the Rules object interface.
Reference: https://aws.amazon.com/whitepapers/security-at-scaIe-governance-in-aws/

NEW QUESTION 5
How does Amazon RDS multi Availability Zone model work?

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

Answer: A

Explanation:
In a Multi-AZ deployment, Amazon RDS automatically provisions and maintains a synchronous standby replica in a different Availability Zone.

Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.Mu|tiAZ.html

NEW QUESTION 6
Which of these is not an instrinsic function in AWS CloudFormation?

A. Fn::EquaIs
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::EquaIs, Fn::If, Fn::Not, Fn::Or, Fn::FindInMap, Fn::GetAtt, Fn::GetAZs, Fn::Join, Fn::Se|ect, Ref
Reference:
http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference.html

NEW QUESTION 7
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/DeveIoperGuide/AutoScaIingGroupLifecycIe.html

NEW QUESTION 8
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. CIoudWatch 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/DeveIoperGuide/dns-failover-determining-health-of-endpoint s.html

NEW QUESTION 9
You run accounting software in the AWS cloud. This software needs to be online continuously during the day every day of the week, and has a very static requirement for compute resources. You also have other, unrelated batch jobs that need to run once per day at any time of your choosing. How should you minimize cost?

A. Purchase a Heavy Utilization Reserved Instance to run the accounting softwar
B. Turn it off after hour
C. Run the batch jobs with the same instance class, so the Reserved Instance credits are also applied to the batch jobs.
D. Purchase a Medium Utilization Reserved Instance to run the accounting softwar
E. Turn it off after hour
F. Run the batch jobs with the same instance class, so the Reserved Instance credits are also applied to the batch jobs.
G. Purchase a Light Utilization Reserved Instance to run the accounting softwar
H. Turn it off after hour
I. Run the batch jobs with the same instance class, so the Reserved Instance credits are also applied to the batch jobs.
J. Purchase a Full Utilization Reserved Instance to run the accounting softwar
K. Turn it off after hour
L. Run the batch jobs with the same instance class, so the Reserved Instance credits are also applied to the batch jobs.

Answer: A

Explanation:
Because the instance will always be online during the day, in a predictable manner, and there are a sequence of batch jobs to perform at any time, we should run the batch jobs when the account software is off. We can achieve Heavy Utilization by alternating these times, so we should purchase the reservation as such, as this represents the lowest cost. There is no such thing a "FuII" level utilization purchases on EC2.
Reference: https://d0.awsstatic.com/whitepapers/Cost_Optimization_with_AWS.pdf

NEW QUESTION 10

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
You are getting a lot of empty receive requests when using Amazon SQS. This is making a lot of unnecessary network load on your instances. What can you do to reduce this load?

A. Subscribe your queue to an SNS topic instead.
B. Use as long of a poll as possible, instead of short polls.
C. Alter your visibility timeout to be shorter.
D. Use <code>sqsd</code> on your EC2 instance

Answer: B

Explanation:
One benefit of long polling with Amazon SQS is the reduction of the number of empty responses, when there are no messages available to return, in reply to a ReceiveMessage request sent to an Amazon SQS queue. Long polling allows the Amazon SQS service to wait until a message is available in the queue before sending a response.
Reference:
http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-long-polling.html

NEW QUESTION 15
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
What is the maximum supported single-volume throughput on EBS?

A. 320IV|iB/s
B. 160MiB/s
C. 40MiB/s
D. 640MiB/s

Answer: A

Explanation:

The ceiling throughput for PIOPS on EBS is 320MiB/s.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVo|umeTypes.htm| IIIIIEZIIII HWS-IIEIIII|]S-EII§iII|}|}I'-PI'0I|}SSi0IIiII EIIEIII

NEW QUESTION 24
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/DeveIoperGuide/AutoScalingGroupLifecycIe.html

NEW QUESTION 28
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-aIarms.htmI

NEW QUESTION 32
Why are more frequent snapshots or EBS Volumes faster?

A. Blocks in EBS Volumes are allocated lazily, since while logically separated from other EBS Volumes, Volumes often share the same physical hardwar
B. Snapshotting the first time forces full block range allocation, so the second snapshot doesn't need to perform the allocation phase and is faster.
C. The snapshots are incremental so that only the blocks on the device that have changed after your last snapshot are saved in the new snapshot.
D. AWS provisions more disk throughput for burst capacity during snapshots if the drive has been pre-warmed by snapshotting and reading all blocks.
E. The drive is pre-warmed, so block access is more rapid for volumes when every block on the device has already been read at least one time.

Answer: B

Explanation:
After writing data to an EBS volume, you can periodically create a snapshot of the volume to use as a baseline for new volumes or for data backup. If you make periodic snapshots of a volume, the snapshots are incremental so that only the blocks on the device that have changed after your last snapshot are saved in the new snapshot. Even though snapshots are saved incrementally, the snapshot deletion process is designed so that you need to retain only the most recent snapshot in order to restore the volume.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-creating-snapshot.html

NEW QUESTION 35
You meet once per month with your operations team to review the past month's data. During the meeting, you realize that 3 weeks ago, your monitoring system which pings over HTTP from outside AWS recorded a large spike in latency on your 3-tier web service API.
You use DynamoDB for the database layer, ELB, EBS, and EC2 for the business logic tier, and SQS, ELB, and EC2 for the presentation layer.
Which of the following techniques will NOT help you figure out what happened?

A. Check your CIoudTraiI log history around the spike's time for any API calls that caused slowness.
B. Review CIoudWatch Metrics graphs to determine which component(s) slowed the system down.
C. Review your ELB access logs in S3 to see if any ELBs in your system saw the latency.
D. Analyze your logs to detect bursts in traffic at that tim

Answer: B

Explanation:
Metrics data are available for 2 weeks. If you want to store metrics data beyond that duration, you can retrieve it using our GetMetricStatistics API as well as a number of applications and tools offered by AWS partners.
Reference: https://aws.amazon.com/cIoudwatch/faqs/

NEW QUESTION 39
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 40
Your team wants to begin practicing continuous delivery using CloudFormation, to enable automated builds and deploys of whole, versioned stacks or stack layers. You have a 3-tier, mission-critical system. Which of the following is NOT a best practice for using CloudFormation in a continuous delivery environment?

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

Answer: B

Explanation:
Putting all resources in one stack is a bad idea, since different tiers have different life cycles and frequencies of change. For additional guidance about organizing your stacks, you can use two common frameworks: a multi-layered architecture and service-oriented architecture (SOA).
Reference:
http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html#organizingstack

NEW QUESTION 41
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://bIogs.aws.amazon.com/security/post/Tx3SYCORF5EKRCO/How-Does-Amazon-Cognito-Relate-to
-Existing-Web-Identity-Federatio

NEW QUESTION 42
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.
"UnavaiIabIe" 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 44
Which of these configuration or deployment practices is a security risk for RDS?

A. Storing SQL function code in plaintext
B. Non-MuIti-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.html

NEW QUESTION 47
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 MuIti-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 49
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 newjobs 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 correctl

Answer: A

Explanation:
The IAM Role must be fine, as if it were broken, NO jobs would be processed since the system would never be able to get any queue messages. The same reasoning applies to the routing table change. The scaling metric is fine, as instance count increased when the queue depth increased due to more messages entering than exiting. Thus, the only reasonable option is that some of the recent messages must be malformed and unprocessable.
Reference:
https://github.com/andrew-templeton/cloudacademy/blob/fca920b45234bbe99cc0e8efb9c65134884dd48 9/questions/null

NEW QUESTION 52
There is a very serious outage at AWS. EC2 is not affected, but your EC2 instance deployment scripts stopped working in the region with the outage. What might be the issue?

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

Answer: B

Explanation:
S3 stores all snapshots. If S3 is unavailable, snapshots are unavailable.
Amazon EC2 also uses Amazon S3 to store snapshots (backup copies) of the data volumes. You can use snapshots for recovering data quickly and reliably in case of application or system failures. You can also use snapshots as a baseline to create multiple new data volumes, expand the size of an existing data volume, or move data volumes across multiple Availability Zones, thereby making your data usage highly scalable. For more information about using data volumes and snapshots, see Amazon Elastic Block Store.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonS3.htmI

NEW QUESTION 53
What is a circular dependency in AWS CIoudFormation?

A. When a Template references an earlier version of itself.
B. When Nested Stacks depend on each other.
C. When Resources form a DependOn loop.
D. When a Template references a region, which references the original Templat

Answer: C

Explanation:
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 CIoudFormation 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 additional information, see DependsOn Attribute.
Reference: http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/troubleshooting.htm|#troub|eshootin g-errors-dependence-error

NEW QUESTION 58
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 CIoudFormation 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 CIoudFormation 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/Iambda/latest/dg/current-supported-versions.html

NEW QUESTION 60
You are creating a new API for video game scores. Reads are 100 times more common than writes, and the top 1% of scores are read 100 times more frequently than the rest of the scores. What's the best design for this system, using DynamoDB?

A. DynamoDB table with 100x higher read than write throughput, with CloudFront caching.
B. DynamoDB table with roughly equal read and write throughput, with CloudFront caching.
C. DynamoDB table with 100x higher read than write throughput, with E|astiCache caching.
D. DynamoDB table with roughly equal read and write throughput, with EIastiCache cachin

Answer: D

Explanation:
Because the 100x read ratio is mostly driven by a small subset, with caching, only a roughly equal number of reads to writes will miss the cache, since the supermajority will hit the top 1% scores. Knowing we need to set the values roughly equal when using caching, we select AWS EIastiCache, because CIoudFront cannot directly cache DynamoDB queries, and EIastiCache is an excellent in-memory cache for database queries, rather than a distributed proxy cache for content delivery.
One solution would be to cache these reads at the application layer. Caching is a technique that is used in many high-throughput applications, offloading read actMty on hot items to the cache rather than to the database. Your application can cache the most popular items in memory, or use a product such as EIastiCache to do the same.
Reference: http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GuideIinesForTabIes.htmI#GuideIi nesForTabIes.CachePopuIarItem

NEW QUESTION 64
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 65

.....

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