



Amazon

Exam Questions AWS-Certified-DevOps-Engineer-Professional

Amazon AWS Certified DevOps Engineer Professional

NEW QUESTION 1

What method should I use to author automation if I want to wait for a CloudFormation stack to finish completing in a script?

- A. Event subscription using SQS.
- B. Event subscription using SNS.
- C. Poll using `ListStacks` / `list-stacks`.
- D. Poll using `GetStackStatus` / `get-stack-status`.

Answer: C

Explanation:

Event driven systems are good for IFTTT logic, but only polling will make a script wait to complete. ListStacks / list-stacks is a real method, GetStackStatus / get-stack-status is not.

Reference: <http://docs.aws.amazon.com/cli/latest/reference/cloudformation/list-stacks.html>

NEW QUESTION 2

You need to perform ad-hoc business analytics queries on well-structured data. Data comes in constantly at a high velocity. Your business intelligence team can understand SQL. What AWS service(s) should you look to first?

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

Answer: B

Explanation:

Kinesis Firehose provides a managed service for aggregating streaming data and inserting it into RedShift. RedShift also supports ad-hoc queries over well-structured data using a SQL-compliant wire protocol, so the business team should be able to adopt this system easily.

Reference: <https://aws.amazon.com/kinesis/firehose/details/>

NEW QUESTION 3

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

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

Answer: A

Explanation:

Terminating SSL terminates the security of a connection over HTTP, removing the S for "Secure" in HTTPS. This violates the "encryption in transit" requirement in the scenario.

Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/elb-listener-config.html>

NEW QUESTION 4

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 5

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

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

- J. Begin using CloudWatch Logs on every service
- K. Stream all Log Groups into an AWS Elasticsearch Service Domain running Kibana 4 and perform log analysis on a search cluster.

Answer: D

Explanation:

The Elasticsearch and Kibana 4 combination is called the ELK Stack, and is designed specifically for real-time, ad-hoc log analysis and aggregation. All other answers introduce extra delay or require pre-defined queries. Amazon Elasticsearch Service is a managed service that makes it easy to deploy, operate, and scale Elasticsearch in the AWS Cloud. Elasticsearch is a popular open-source search and analytics engine for use cases such as log analytics, real-time application monitoring, and click stream analytics. Reference: <https://aws.amazon.com/elasticsearch-service/>

NEW QUESTION 6

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

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

Answer: A

Explanation:

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-endpoints.html>

NEW QUESTION 7

What is the scope of an EC2 security group?

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

Answer: C

Explanation:

A security group is tied to a region and can be assigned only to instances in the same region. You can't enable an instance to communicate with an instance outside its region using security group rules. Traffic from an instance in another region is seen as WAN bandwidth. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/resources.html>

NEW QUESTION 8

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

- A. Use Kinesis as a buffer stream for message bodies
- 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 medium
- 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 9

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 `qsqs` 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 10

You run operations for a company that processes digital wallet payments at a very high volume. One second of downtime, during which you drop payments or are otherwise unavailable, loses you on average USD 100. You balance the financials of the transaction system once per day. Which database setup is best suited to address this business risk?

- A. A multi-AZ RDS deployment with synchronous replication to multiple standbys and read-replicas for fast failover and ACID properties.
- B. A multi-region, multi-master, active-active RDS configuration using database-level ACID design principles with database trigger writes for replication.
- C. A multi-region, multi-master, active-active DynamoDB configuration using application control-level BASE design principles with change-stream write queue buffers for replication.
- D. A multi-AZ DynamoDB setup with changes streamed to S3 via AWS Kinesis, for highly durable storage and BASE properties.

Answer: C

Explanation:

Only the multi-master, multi-region DynamoDB answer makes sense. Multi-AZ deployments do not provide sufficient availability when a business loses USD 360,000 per hour of unavailability. As RDS does not natively support multi-region, and ACID does not perform well/at all over large distances between regions, only the DynamoDB answer works. Reference:

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.CrossRegionRepl.html>

NEW QUESTION 10

Which of the following are not valid sources for OpsWorks custom cookbook repositories?

- A. HTTP(S)
- B. Git
- C. AWS EBS
- D. Subversion

Answer: C

Explanation:

Linux stacks can install custom cookbooks from any of the following repository types: HTTP or Amazon S3 archives. They can be either public or private, but Amazon S3 is typically the preferred option for a private archive. Git and Subversion repositories provide source control and the ability to have multiple versions. Reference:

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

NEW QUESTION 15

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

- A. 320 MiB/s
- B. 160 MiB/s
- C. 40 MiB/s
- D. 640 MiB/s

Answer: A

Explanation:

The ceiling throughput for PIOPS on EBS is 320 MiB/s.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html>

NEW QUESTION 20

Your serverless architecture using AWS API Gateway, AWS Lambda, and AWS DynamoDB experienced a large increase in traffic to a sustained 400 requests per second, and dramatically increased in failure rates. Your requests, during normal operation, last 500 milliseconds on average. Your DynamoDB table did not exceed 50% of provisioned throughput, and Table primary keys are designed correctly. What is the most likely issue?

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

Answer: C

Explanation:

AWS API Gateway by default throttles at 500 requests per second steady-state, and 1000 requests per second at spike. Lambda, by default, throttles at 100 concurrent requests for safety. At 500 milliseconds (half of a second) per request, you can expect to support 200 requests per second at 100 concurrency. This is less than the 400 requests per second your system now requires. Make a limit increase request via the AWS Support Console.

AWS Lambda: Concurrent requests safety throttle per account -> 100

Reference: http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#limits_lambda

NEW QUESTION 24

You need to migrate 10 million records in one hour into DynamoDB. All records are 1.5KB in size. The data is evenly distributed across the partition key. How many write capacity units should you provision during this batch load?

- A. 6667
- B. 4166

- C. 5556
- D. 2778

Answer: C

Explanation:

You need 2 units to make a 1.5KB write, since you round up. You need 20 million total units to perform this load. You have 3600 seconds to do so. DMde and round up for 5556.

Reference: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ProvisionedThroughput.html>

NEW QUESTION 28

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 30

You are experiencing performance issues writing to a DynamoDB table. Your system tracks high scores for video games on a marketplace. Your most popular game experiences all of the performance issues. What is the most likely problem?

- A. DynamoDB's vector clock is out of sync, because of the rapid growth in request for the most popular game.
- B. You selected the Game ID or equivalent identifier as the primary partition key for the table.
- C. Users of the most popular video game each perform more read and write requests than average.
- D. You did not provision enough read or write throughput to the table.

Answer: B

Explanation:

The primary key selection dramatically affects performance consistency when reading or writing to DynamoDB. By selecting a key that is tied to the identity of the game, you forced DynamoDB to create a hotspot in the table partitions, and over-request against the primary key partition for the popular game. When it stores data, DynamoDB divides a table's items into multiple partitions, and distributes the data primarily based upon the partition key value. The provisioned throughput associated with a table is also divided evenly among the partitions, with no sharing of provisioned throughput across partitions. Reference:

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GuidelinesForTables.html#GuidelinesForTables.UniformWorkload>

NEW QUESTION 34

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 a 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 Microsoft 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 35

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 37

Which of these is not a CloudFormation Helper Script?

- A. cfn-signal
- B. cfn-hup
- C. cfn-request
- D. cfn-get-metadata

Answer: C

Explanation:

This is the complete list of CloudFormation Helper Scripts: cfn-init, cfn-signal, cfn-get-metadata, cfn-hup Reference:
<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-helper-scripts-reference.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 `ValidateTemplate` call before publishing changes to AWS.
- B. Model your stack in one template, so you can leverage CloudFormation's state management and dependency resolution to propagate all changes.
- C. Use CloudFormation to create brand new infrastructure for all stateless resources on each push, and run integration tests on that set of infrastructure.
- D. Parametrize the template and use `Mappings` to ensure your template works in multiple Regions.

Answer: B

Explanation:

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://blogs.aws.amazon.com/security/post/Tx3SYCORF5EKRCO/How-Does-Amazon-Cognito-Relate-to-Existing-Web-Identity-Federatio>

NEW QUESTION 43

You need to deploy a new application version to production. Because the deployment is high-risk, you need to roll the new version out to users over a number of hours, to make sure everything is working correctly. You need to be able to control the proportion of users seeing the new version of the application down to the percentage point.

You use ELB and EC2 with Auto Scaling Groups and custom AMIs with your code pre-installed assigned to Launch Configurations. There are no database-level changes during your deployment. You have been told you cannot spend too much money, so you must not increase the number of EC2 instances much at all during the deployment, but you also need to be able to switch back to the original version of code quickly if something goes wrong. What is the best way to meet these requirements?

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

Answer: A

Explanation:

Only Weighted Round Robin DNS Records and reverse proxies allow such fine-grained tuning of traffic splits. The Blue-Green option does not meet the requirement that we mitigate costs and keep overall EC2 fleet size consistent, so we must select the 2 ELB and ASG option with WRR DNS tuning. This method is called A/B deployment and/or Canary deployment.

Reference: <https://d0.awsstatic.com/whitepapers/overview-of-deployment-options-on-aws.pdf>

NEW QUESTION 48

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.html>

NEW QUESTION 53

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 region
- B. Create an Auto Scaling Group behind an ELB in each of the two regions DynamoDB is running in
- 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 table
- E. Create an Auto Scaling Group behind an ELB in each of the two regions DynamoDB is running in
- 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 Multi-Region table
- 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 region
- 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.CrossRegionRepl.html>

NEW QUESTION 55

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 57

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 architecture for this system?

- A. Create an OpsWorks Stack with two Layers
- 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 files
- C. When a new file is detected, request instances to boot on the ANN layer
- D. When the instances are booted and the HTTP APIs are up, submit processing requests to individual instances.

- E. Make an S3 notification configuration which publishes to AWS Lambda on the manifest bucket
- 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 instance
- G. Create an SQS queue of the images in the manifest
- 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++ extension
- J. Make an S3 notification configuration on the manifest, which publishes to another AWS Lambda running controller code
- K. This controller code publishes all the images in the manifest to AWS Kinesis
- L. Your ANN code Lambda Function uses the Kinesis as an Event Source
- M. The system automatically scales when the stream contains image events.
- N. Create an Auto Scaling, Load Balanced Elastic Beanstalk worker tier Application and Environment
- O. Deploy the ANN code to G2 instances in this tier
- 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 fleet 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 59

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 ElastiCache caching.
- D. DynamoDB table with roughly equal read and write throughput, with ElastiCache caching.

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 ElastiCache, because CloudFront cannot directly cache DynamoDB queries, and ElastiCache 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 activity 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 ElastiCache to do the same.

Reference: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GuidelinesForTables.html#GuidelinesForTables.CachePopularItem>

NEW QUESTION 60

You need the absolute highest possible network performance for a cluster computing application. You already selected homogeneous instance types supporting 10 gigabit enhanced networking, made sure that your workload was network bound, and put the instances in a placement group. What is the last optimization you can make?

- A. Use 9001 MTU instead of 1500 for Jumbo Frames, to raise packet body to packet overhead ratios.
- B. Segregate the instances into different peered VPCs while keeping them all in a placement group, so each one has its own Internet Gateway.
- C. Bake an AMI for the instances and relaunch, so the instances are fresh in the placement group and don't have noisy neighbors.
- D. Turn off SYN/ACK on your TCP stack or begin using UDP for higher throughput

Answer: A

Explanation:

For instances that are colocated inside a placement group, jumbo frames help to achieve the maximum network throughput possible, and they are recommended in this case. For more information, see Placement Groups.

Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/network_mtu.html#jumbo_frame_instances

NEW QUESTION 62

You have a high security requirement for your AWS accounts. What is the most rapid and sophisticated setup you can use to react to AWS API calls to your account?

- A. Subscription to AWS Config via an SNS Topic
- B. Use a Lambda Function to perform in-flight analysis and react to changes as they occur.
- C. Global AWS CloudTrail setup delivering to S3 with an SNS subscription to the deliver notifications, pushing into a Lambda, which inserts records into an ELK stack for analysis.
- D. Use a CloudWatch Rule ScheduleExpression to periodically analyze IAM credential log
- E. Push the deltas for events into an ELK stack and perform ad-hoc analysis there.
- F. CloudWatch Events Rules which trigger based on all AWS API calls, submitting all events to an AWS Kinesis Stream for arbitrary downstream analysis.

Answer: D

Explanation:

CloudWatch Events allow subscription to AWS API calls, and direction of these events into Kinesis Streams. This allows a unified, near real-time stream for all API calls, which can be analyzed with any tool(s) of your choosing downstream.

Reference: http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/EventTypes.html#api_event_type

NEW QUESTION 64

.....

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

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](#)