AWS-Certified-DevOps-Engineer-Professional Dumps

Amazon AWS Certified DevOps Engineer Professional

https://www.certleader.com/AWS-Certified-DevOps-Engineer-Professional-dumps.html

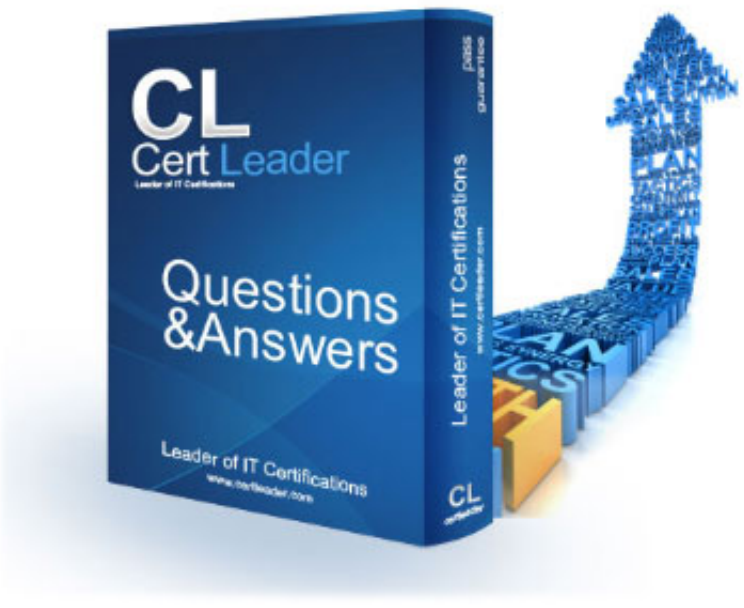

NEW QUESTION 1
You need to perform ad-hoc business analytics queries on well-structured data. Data comes in constantly at a high velocity. Your business intelligence team can understand SQL. What AWS service(s) should you look to first?

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/detai|s/

NEW QUESTION 2
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/GuideIinesForTabIes.htmI#GuideIi nesForTabIes.Partitions

NEW QUESTION 3
What is server immutability?

A. Not updating a server after creation.
B. The ability to change server counts.
C. Updating a server after creation.
D. The inability to change server count

Answer: A

Explanation: disposable upgrades offer a simpler way to know if your application has unknown dependencies. The underlying EC2 instance usage is considered temporary or ephemeral in nature for the period of deployment until the current release is active. During the new release, a new set of EC2 instances are rolled out by terminating older instances. This type of upgrade technique is more common in an immutable infrastructure.
Reference: https://d0.awsstatic.com/whitepapers/overview-of-deployment-options-on-aws.pdf

NEW QUESTION 4
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 5
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 CIoudFront.
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 6
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 CIoudFormation, 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 7
When thinking of DynamoDB, what are true of Global Secondary Key properties?

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

Answer: A

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 "gIobaI" 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 8
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/MessageLifecycIe.ht ml

NEW QUESTION 9
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 10
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 MuIti-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.htmI

NEW QUESTION 10
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 instanc
B. Use an S3 Notification Configuration on the <code>PutBucket</code> event and publish events to AWS Lambd
C. Use the Lambda to analyze logs as soon as they come in and flag issues.
D. Begin using CIoudWatch Logs on every servic
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 instanc
H. Use an S3 Notification Configuration on the <code>PutBucket</code> event and publish events to AWS Kinesi
I. Use Apache Spark on AWS EMR to perform at-scale stream processing queries on the log chunks and flag issues.
J. Begin using CIoudWatch Logs on every servic
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 12
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.htmI

NEW QUESTION 14
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.htmI

NEW QUESTION 18
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 22
Which EBS volume type is best for high performance NoSQL cluster deployments?

A. iol
B. gpl
C. standard
D. gp2

Answer: A

Explanation: io1 volumes, or Provisioned IOPS (PIOPS) SSDs, are best for: Critical business applications that require sustained IOPS performance, or more than 10,000 IOPS or 160 MiB/s of throughput per volume, like large database workloads, such as MongoDB.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVo|umeTypes.htm|

NEW QUESTION 27
You are building out a layer in a software stack on AWS that needs to be able to scale out to react to increased demand as fast as possible. You are running the code on EC2 instances in an Auto Scaling Group behind an ELB. Which application code deployment method should you use?

A. SSH into new instances that come online, and deploy new code onto the system by pulling it from an S3 bucket, which is populated by code that you refresh from source control on new pushes.
B. Bake an AMI when deploying new versions of code, and use that AMI for the Auto Scaling Launch Configuration.
C. Create a Dockerfile when preparing to deploy a new version to production and publish it to S3. Use UserData in the Auto Scaling Launch configuration to pull down the Dockerfile from S3 and run it when new instances launch.
D. Create a new Auto Scaling Launch Configuration with UserData scripts configured to pull the latest code at all times.

Answer: B

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

NEW QUESTION 31
You need to perform ad-hoc analysis on log data, including searching quickly for specific error codes and reference numbers. Which should you evaluate first?

A. AWS Elasticsearch Service
B. AWS RedShift
C. AWS EMR
D. AWS DynamoDB

Answer: A

Explanation: Amazon Elasticsearch Service (Amazon ES) is a managed service that makes it easy to deploy, operate, and scale Elasticsearch clusters 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:
http://docs.aws.amazon.com/elasticsearch-service/latest/developerguide/what-is-amazon-elasticsearch-s ervice.html

NEW QUESTION 32
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 37
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 40
When thinking of AWS Elastic BeanstaIk's model, which is true?

A. Applications have many deployments, deployments have many environments.
B. Environments have many applications, applications have many deployments.
C. Applications have many environments, environments have many deployments.
D. Deployments have many environments, environments have many application

Answer: C

Explanation: Applications group logical services. Environments belong to Applications, and typically represent different deployment levels (dev, stage, prod, fo forth). Deployments belong to environments, and are pushes of bundles of code for the environments to run.
Reference: http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/\NeIcome.html

NEW QUESTION 41
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 44
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 48
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. IV|u|ti-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.CrossRegionRepI.html

NEW QUESTION 49
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 "gIobaI" 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.htmI

NEW QUESTION 53
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 57
Which major database needs a BYO license?

A. PostgreSQL
B. NIariaDB
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_OracIe.htm|

NEW QUESTION 62
You need to grant a vendor access to your AWS account. They need to be able to read protected messages in a private S3 bucket at their leisure. They also use AWS. What is the best way to accomplish this?

A. Create an IAM User with API Access Key
B. Grant the User permissions to access the bucke
C. Give the vendor the AWS Access Key ID and AWS Secret Access Key for the User.
D. Create an EC2 Instance Profile on your accoun
E. Grant the associated IAM role full access to the bucke
F. Start an EC2 instance with this Profile and give SSH access to the instance to the vendor.
G. Create a cross-account IAM Role with permission to access the bucket, and grant permission to use the Role to the vendor AWS account.
H. Generate a signed S3 PUT URL and a signed S3 PUT URL, both with wildcard values and 2 year duration
I. Pass the URLs to the vendor.

Answer: C

Explanation: When third parties require access to your organization's AWS resources, you can use roles to delegate access to them. For example, a third party might provide a service for managing your AWS resources. With IAM roles, you can grant these third parties access to your AWS resources without sharing your AWS security credentials. Instead, the third party can access your AWS resources by assuming a role that you create in your AWS account.
Reference:
http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roIes_common-scenarios_third-party.html

NEW QUESTION 63
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 65
For AWS CloudFormation, which stack state refuses UpdateStack calls?

A. <code>UPDATE_ROLLBACK_FAILED</code>
B. <code>UPDATE_ROLLBACK_COMPLETE</code>
C. <code>UPDATE_CONIPLETE</code>
D. <code>CREATE_COMPLETE</code>

Answer: A

Explanation: When a stack is in the UPDATE_ROLLBACK_FA|LED state, you can continue rolling it back to return it to a working state (to UPDATE_ROLLBACK_COMPLETE). You cannot update a stack that is in the UPDATE_ROLLBACK_FA|LED state. However, if you can continue to roll it back, you can return the stack to its original settings and try to update it again.
Reference:
http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/using-cfn-updating-stacks-continueu pdateroIIback.html

NEW QUESTION 66
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.ProvisionedThroughp ut.html

NEW QUESTION 71
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 76
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 tabl

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 dMdes a tabIe'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 dMded evenly among the partitions, with no sharing of provisioned throughput across partitions. Reference: http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GuideIinesForTabIes.html#GuideIi nesForTabIes.UniformWorkIoad

NEW QUESTION 78
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 81
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 CIoudFront 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/EBSVoIumeTypes.htmI

NEW QUESTION 84
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. CIoudFormation 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 CIoudFormation 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 CIoudFormation, 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 CIoudFormation 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 CIoudFormation template and have AWS CIoudFormation 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/AWSCIoudFormation/latest/UserGuide/template-custom-resources.html

NEW QUESTION 89
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/detaiIs/

NEW QUESTION 94
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 96
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/AWSCIoudFormation/latest/UserGuide/best-practices.html#organizingstack

NEW QUESTION 97
You need to replicate API calls across two systems in real time. What tool should you use as a buffer and transport mechanism for API call events?

A. AWS SQS
B. AWS Lambda
C. AWS Kinesis
D. AWS SNS

Answer: C

Explanation: AWS Kinesis is an event stream service. Streams can act as buffers and transport across systems for in-order programmatic events, making it ideal for replicating API calls across systems.
A typical Amazon Kinesis Streams application reads data from an Amazon Kinesis stream as data records. These applications can use the Amazon Kinesis Client Library, and they can run on Amazon EC2 instances. The processed records can be sent to dashboards, used to generate alerts, dynamically change pricing and advertising strategies, or send data to a variety of other AWS services. For information about Streams features and pricing, see Amazon Kinesis Streams.
Reference: http://docs.aws.amazon.com/kinesis/Iatest/dev/introduction.html

NEW QUESTION 100
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 BeanstaIk'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/elasticbeanstaIk/Iatest/dg/create_deploy_Ruby_raiIs.html

NEW QUESTION 101
What is the scope of AWS IAM?

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

Answer: A

Explanation: IAM resources are all global; there is not regional constraint. Reference: https://aws.amazon.com/iam/faqs/

NEW QUESTION 106
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 110
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 CloudTraiI 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 CloudTraiI.
CloudTrai| provides visibility into user actMty by recording API calls made on your account. CloudTraiI 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. CloudTraiI makes it easier to ensure compliance with internal policies and regulatory standards. Reference: https://aws.amazon.com/CloudtraiI/faqs/

NEW QUESTION 112
What is the order of most-to-least rapidly-scaling (fastest to scale first)?
(A) EC2 + ELB + Auto Scaling (B) Lambda (C) RDS

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

Answer: A

Explanation: Lambda is designed to scale instantly. EC2 + ELB + Auto Scaling require single-digit minutes to scale out. RDS will take atleast 15 minutes, and will apply OS patches or any other updates when applied. Reference: https://aws.amazon.com/|ambda/faqs/

NEW QUESTION 114
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 117
What is required to achieve gigabit network throughput on EC2? You already selected cluster-compute, 10GB instances with enhanced networking, and your workload is already network-bound, but you are not seeing 10 gigabit speeds.

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

Answer: D

Explanation: You are not guaranteed 10gigabit performance, except within a placement group.
A placement group is a logical grouping of instances within a single Availability Zone. Using placement groups enables applications to participate in a low-latency, 10 Gbps network. Placement groups are recommended for applications that benefit from low network latency, high network throughput, or both. Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html

NEW QUESTION 119
If I want CIoudFormation 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 CIoudFormation 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 CIoudFormation 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 CIoudFormation stack to publish events into.
D. Subscribe your continuous delivery system to an SQS queue that you also tell your CIoudFormation 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/AWSCIoudFormation/latest/UserGuide/using-cfn-updating-stacks-monitor-s tack.html

NEW QUESTION 121
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 124
From a compliance and security perspective, which of these statements is true?

A. You do not ever need to rotate access keys for AWS IAM Users.
B. You do not ever need to rotate access keys for AWS IAM Roles, nor AWS IAM Users.
C. None of the other statements are true.
D. You do not ever need to rotate access keys for AWS IAM Role

Answer: D

Explanation: IAM Role Access Keys are auto-rotated by AWS on your behalf; you do not need to rotate them.
The application is granted the permissions for the actions and resources that you've defined for the role through the security credentials associated with the role. These security credentials are temporary and we
rotate them automatically. We make new credentials available at least five minutes prior to the expiration of the old credentials.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html

NEW QUESTION 125
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 129
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.html

NEW QUESTION 131
Which of the following tools does not directly support AWS OpsWorks, for monitoring your stacks?

A. AWS Config
B. Amazon CIoudWatch Nletrics
C. AWS CloudTraiI
D. Amazon CIoudWatch Logs

Answer: A

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

NEW QUESTION 135
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 139
You need to run a very large batch data processing job one time per day. The source data exists entirely in S3, and the output of the processing job should also be written to S3 when finished. If you need to version control this processing job and all setup and teardown logic for the system, what approach should you use?

A. Model an AWS EMR job in AWS Elastic Beanstalk.
B. Model an AWS EMR job in AWS CloudFormation.
C. Model an AWS EMR job in AWS OpsWorks.
D. Model an AWS EMR job in AWS CLI Compose

Answer: B

Explanation: To declaratively model build and destroy of a cluster, you need to use AWS CIoudFormation. OpsWorks and Elastic Beanstalk cannot directly model EMR Clusters. The CLI is not declarative, and CLI Composer does not exist.
Reference:
http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/aws-resource-emr-cluster.html

NEW QUESTION 141
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/weIcome.html

NEW QUESTION 146
You are designing a system which needs, at minimum, 8 m4.Iarge instances operating to service traffic. When designing a system for high availability in the us-east-1 region, which has 6 Availability Zones, you company needs to be able to handle death of a full availability zone. How should you distribute the servers, to save as much cost as possible, assuming all of the EC2 nodes are properly linked to an ELB? Your VPC account can utilize us-east-1's AZ's a through f, inclusive.

A. 3 servers in each of AZ's a through d, inclusive.
B. 8 servers in each of AZ's a and b.
C. 2 servers in each of AZ's a through e, inclusive.
D. 4 servers in each of AZ's a through c, inclusiv

Answer: C

Explanation: You need to design for N+1 redundancy on Availability Zones. ZONE_COUNT = (REQUIRED_INSTANCES / INSTANCE_COUNT_PER_ZONE) + 1. To minimize cost, spread the instances across as many possible zones as you can. By using a though e, you are allocating 5 zones. Using 2 instances, you have 10 total instances. If a single zone fails, you have 4 zones left, with 2 instances each, for a total of 8 instances. By spreading out as much as possible, you have increased cost by only 25% and significantly de-risked an availability zone failure.
Reference:
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html#concepts- regions-availability-zones

NEW QUESTION 149
You need to create a Route53 record automatically in CIoudFormation 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 150
What is web identity federation?

A. Use of an identity provider like Google or Facebook to become an AWS IAM User.
B. Use of an identity provider like Google or Facebook to exchange for temporary AWS security credentials.
C. Use of AWS IAM User tokens to log in as a Google or Facebook user.
D. Use of AWS STS Tokens to log in as a Google or Facebook use

Answer: B

Explanation: users of your app can sign in using a well-known identity provider (|dP) -such as Login with Amazon, Facebook, Google, or any other OpenID Connect (OIDC)-compatible IdP, receive an authentication token, and then exchange that token for temporary security credentials in AWS that map to an IAM role with permissions to use the resources in your AWS account.
Reference: http://docs.aws.amazon.com/IANI/latest/UserGuide/id_roIes_providers_oidc.html

NEW QUESTION 151
You have been asked to de-risk deployments at your company. Specifically, the CEO is concerned about outages that occur because of accidental inconsistencies between Staging and Production, which sometimes cause unexpected behaviors in Production even when Staging tests pass.
You already use Docker to get high consistency between Staging and Production for the application environment on your EC2 instances. How do you further de-risk the rest of the execution environment, since in AWS, there are many service components you may use beyond EC2 virtual machines?

A. Develop models of your entire cloud system in CIoudFormatio

B. Use this model in Staging and Production to achieve greater parity.
C. Use AWS Config to force the Staging and Production stacks to have configuration parit
D. Any differences will be detected for you so you are aware of risks.
E. Use AMIs to ensure the whole machine, including the kernel of the virual machines, is consistent, since Docker uses Linux Container (LXC) technology, and we need to make sure the container environment is consistent.
F. Use AWS ECS and Docker clusterin
G. This will make sure that the AMIs and machine sizes are the same across both environments.

Answer: A

Explanation: Only CIoudFormation's JSON Templates allow declarative version control of repeatably deployable models of entire AWS clouds.
Reference: https://bIogs.aws.amazon.com/application-management/blog/category/Best+practices

NEW QUESTION 154
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 156
What is the scope of an EBS snapshot?

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

Answer: C

Explanation: An EBS snapshot is tied to its region and can only be used to create volumes in the same region. You can copy a snapshot from one region to another. For more information, see Copying an Amazon EBS Snapshot.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/resources.html

NEW QUESTION 159
Your system uses a multi-master, multi-region DynamoDB configuration spanning two regions to achieve high availablity. For the first time since launching your system, one of the AWS Regions in which you operate over went down for 3 hours, and the failover worked correctly. However, after recovery, your users are experiencing strange bugs, in which users on different sides of the globe see different data. What is a likely design issue that was not accounted for when launching?

A. The system does not have Lambda Functor Repair Automations, to perform table scans and chack for corrupted partition blocks inside the Table in the recovered Region.
B. The system did not implement DynamoDB Table Defragmentation for restoring partition performance in the Region that experienced an outage, so data is served stale.
C. The system did not include repair logic and request replay buffering logic for post-failure, to re-synchronize data to the Region that was unavailable for a number of hours.
D. The system did not use DynamoDB Consistent Read requests, so the requests in different areas are not utilizing consensus across Regions at runtime.

Answer: C

Explanation: When using multi-region DynamoDB systems, it is of paramount importance to make sure that all requests made to one Region are replicated to the other. Under normal operation, the system in question would correctly perform write replays into the other Region. If a whole Region went down, the system would be unable to perform these writes for the period of downtime. Without buffering write requests somehow, there would be no way for the system to replay dropped cross-region writes, and the requests would be serviced differently depending on the Region from which they were served after recovery. Reference: http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.CrossRegionRepI.htmI

NEW QUESTION 162
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.html

NEW QUESTION 163
Your company needs to automate 3 layers of a large cloud deployment. You want to be able to track this deployment's evolution as it changes over time, and carefully control any alterations. What is a good way to automate a stack to meet these requirements?

A. Use OpsWorks Stacks with three layers to model the layering in your stack.
B. Use CloudFormation Nested Stack Templates, with three child stacks to represent the three logicallayers of your cloud.
C. Use AWS Config to declare a configuration set that AWS should roll out to your cloud.
D. Use Elastic Beanstalk Linked Applications, passing the important DNS entires between layers using the metadata interface.

Answer: B

Explanation: Only CloudFormation allows source controlled, declarative templates as the basis for stack automation. Nested Stacks help achieve clean separation of layers while simultaneously providing a method to control all layers at once when needed.
Reference:
https://blogs.aws.amazon.com/application-management/post/TxIT9JYOOS8AB9I/Use-Nested-Stacks-to- Create-Reusable-Templates-and-Support-Role-Specialization

NEW QUESTION 166
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 donot have noisy neighbors.
D. Turn off SYN/ACK on your TCP stack or begin using UDP for higher throughpu

Answer: A

Explanation: For instances that are collocated 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.htm|#jumbo_frame_instances

NEW QUESTION 171
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_credentialIs_mfa.htmI

NEW QUESTION 176
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/elasticbeanstalk/latest/dg/using-features-managing-env-tiers.htmI#worker-d eadletter

NEW QUESTION 178

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your AWS-Certified-DevOps-Engineer-Professional Exam with Our Prep Materials Via below:

https://www.certleader.com/AWS-Certified-DevOps-Engineer-Professional-dumps.html