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

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

Answer: B

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

NEW QUESTION 2
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 3
You need your CI to build AMIs with code pre-installed on the images on every new code push. You need to do this as cheaply as possible. How do you do this?

A. Bid on spot instances just above the asking price as soon as new commits come in, perform all instance configuration and setup, then create an AMI based on the spot instance.
B. Have the CI launch a new on-demand EC2 instance when new commits come in, perform all instance configuration and setup, then create an AMI based on the on-demand instance.
C. Purchase a Light Utilization Reserved Instance to save money on the continuous integration machin
D. Use these credits whenever your create AMIs on instances.
E. When the CI instance receives commits, attach a new EBS volume to the CI machin
F. Perform all setup on this EBS volume so you don't need a new EC2 instance to create the AMI.

Answer: A

Explanation:
Spot instances are the cheapest option, and you can use minimum run duration if your AMI takes more than a few minutes to create.
Spot instances are also available to run for a predefined duration — in hourly increments up to six hours in length — at a significant discount (30-45%) compared to On-Demand pricing plus an additional 5% during off-peak timesl for a total of up to 50% savings.
Reference: https://aws.amazon.com/ec2/spot/pricing/

NEW QUESTION 4
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 EIPs per Region for VPC on new accounts. The third environment could not allocate the 6th EIP.
Reference: http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#limits_vpc

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

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

Answer:

A

Explanation:
In a Multi-AZ deployment, Amazon RDS automatically provisions and maintains a synchronous standby replica in a different Availability Zone.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.Mu|tiAZ.html

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

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

Answer: D

Explanation:
This is the complete list of Intrinsic Functions...: Fn::Base64, Fn::And, Fn::Equals, Fn::If, Fn::Not, Fn::Or, Fn::FindInMap, Fn::GetAtt, Fn::GetAZs, Fn::Join, Fn::Se|ect, Ref
Reference:
http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference.html

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

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

Answer: A

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

NEW QUESTION 8
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 "Full" level utilization purchases on EC2.
Reference: https://d0.awsstatic.com/whitepapers/Cost_Optimization_with_AWS.pdf

NEW QUESTION 9
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 10
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 10
For AWS Auto Scaling, what is the first transition state an existing instance enters after leaving steady state in Standby mode?

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

Answer: C

Explanation:
You can put any instance that is in an InService state into a Standby state. This enables you to remove the instance from service, troubleshoot or make changes to it, and then put it back into service. Instances in a Standby state continue to be managed by the Auto Scaling group. However, they are not an active part of your application until you put them back into service.
Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AutoScalingGroupLifecycle.html

NEW QUESTION 11
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 13
Which deployment method, when using AWS Auto Scaling Groups and Auto Scaling Launch Configurations, enables the shortest time to live for indMdual sewers?

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

Answer: A

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

NEW QUESTION 16
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 21
Which major database needs a BYO license?

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

Answer: D

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

NEW QUESTION 23
You need to know when you spend $1000 or more on AWS. What's the easy way for you to see that notification?

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

Answer: C

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

NEW QUESTION 28
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 32
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/AWSCloudFormation/latest/UserGuide/using-cfn-updating-stacks-continueu pdaterollback.html

NEW QUESTION 33
Your CTO thinks your AWS account was hacked. What is the only way to know for certain if there was unauthorized access and what they did, assuming your hackers are very sophisticated AWS engineers and doing everything they can to cover their tracks?

A. Use CloudTrai| Log File Integrity Validation.
B. Use AWS Config SNS Subscriptions and process events in real time.
C. Use CloudTrail backed up to AWS S3 and Glacier.
D. Use AWS Config Timeline forensic

Answer: A

Explanation:
You must use CloudTrail Log File Validation (default or custom implementation), as any other tracking method is subject to forgery in the event of a full account compromise by sophisticated enough hackers. Validated log files are invaluable in security and forensic investigations. For example, a validated log file enables you to assert positively that the log file itself has not changed, or that particular user credentials performed specific API actMty. The CloudTrail log file integrity validation process also lets you know if a log file has been deleted or changed, or assert positively that no log files were delivered to your account during a given period of time.
Reference:
http://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-log-fiIe-validation-intro.html

NEW QUESTION 37
Which of these is not a Pseudo Parameter in AWS CIoudFormation?

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

Answer: C

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

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

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

Answer: D

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

NEW QUESTION 42
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 46
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 CIoudTrail log history around the spike's time for any API calls that caused slowness.
B. Review CloudWatch 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/cloudwatch/faqs/

NEW QUESTION 50
Which of these is not an intrinsic function in AWS CloudFormation?

A. Fn::Split
B. Fn::FindInMap
C. Fn::Select
D. Fn::GetAZs

Answer: A

Explanation:
This is the complete list of Intrinsic Functions...: Fn::Base64, Fn::And, Fn::Equals, Fn::If, Fn::Not, Fn::Or, Fn::FindInMap, Fn::GetAtt, Fn::GetAZs, Fn::Join, Fn::Select, Ref
Reference:
http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference.html

NEW QUESTION 54
Your API requires the ability to stay online during AWS regional failures. Your API does not store any state, it only aggregates data from other sources - you do not have a database. What is a simple but effective way to achieve this uptime goal?

A. Use a CloudFront distribution to serve up your AP
B. Even if the region your API is in goes down, the edge locations CloudFront uses will be fine.
C. Use an ELB and a cross-zone ELB deployment to create redundancy across datacenter
D. Even if a region fails, the other AZ will stay online.
E. Create a Route53 Weighted Round Robin record, and if one region goes down, have that region redirect to the other region.
F. Create a Route53 Latency Based Routing Record with Failover and point it to two identical deployments of your stateless API in two different region
G. Make sure both regions use Auto Scaling Groups behind ELBs.

Answer: D

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

NEW QUESTION 55
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 60
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 64
Which is not a restriction on AWS EBS Snapshots?

A. Snapshots which are shared cannot be used as a basis for other snapshots.
B. You cannot share a snapshot containing an AWS Access Key ID or AWS Secret Access Key.
C. You cannot share unencrypted snapshots.

D. Snapshot restorations are restricted to the region in which the snapshots are create

Answer: A

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

NEW QUESTION 69
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 AM|s 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 fileet 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 70
Which of these is not a reason a Mu|ti-AZ RDS instance will failover?

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

Answer: C

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

NEW QUESTION 71
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>DepIoyCode</code> API call when there are major outages, to protect from system floods.
D. None of the other answers make sens
E. If EC2 is not affected, it must be some other issu

Answer: B

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

NEW QUESTION 75

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

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

Answer: A

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

NEW QUESTION 76
You work for a company that automatically tags photographs using artificial neural networks (ANNs), which run on GPUs using C++. You receive millions of images at a time, but only 3 times per day on average. These images are loaded into an AWS S3 bucket you control for you in a batch, and then the customer publishes a JSON-formatted manifest into another S3 bucket you control as well. Each image takes 10 milliseconds to process using a full GPU. Your neural network software requires 5 minutes to bootstrap. Image tags are JSON objects, and you must publish them to an S3 bucket.
Which of these is the best system architectures for this system?

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

Answer: B

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

NEW QUESTION 80
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 84
Your application's Auto Scaling Group scales up too quickly, too much, and stays scaled when traffic decreases. What should you do to fix this?

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

Answer: B

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

NEW QUESTION 89
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_credentials_mfa.html

NEW QUESTION 90
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 Topi
B. Use a Lambda Function to perform in-flight analysis and reactMty 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_ty pe

NEW QUESTION 91
......

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