



Amazon

Exam Questions AWS-Certified-DevOps-Engineer-Professional

Amazon AWS Certified DevOps Engineer Professional

NEW QUESTION 1

Your application consists of 10% writes and 90% reads. You currently service all requests through a Route53 Alias Record directed towards an AWS ELB, which sits in front of an EC2 Auto Scaling Group. Your system is getting very expensive when there are large traffic spikes during certain news events, during which many more people request to read similar data all at the same time. What is the simplest and cheapest way to reduce costs and scale with spikes like this?

- A. Create an S3 bucket and asynchronously replicate common requests responses into S3 object
- B. When a request comes in for a precomputed response, redirect to AWS S3.
- C. Create another ELB and Auto Scaling Group layer mounted on top of the other system, adding a tier to the system
- D. Serve most read requests out of the top layer.
- E. Create a CloudFront Distribution and direct Route53 to the Distribution
- F. Use the ELB as an Origin and specify Cache Behaviours to proxy cache requests which can be served late.
- G. Create a Memcached cluster in AWS ElastiCache
- H. Create cache logic to serve requests which can be served late from the in-memory cache for increased performance.

Answer: C

Explanation:

CloudFront is ideal for scenarios in which entire requests can be served out of a cache and usage patterns involve heavy reads and spikiness in demand. A cache behavior is the set of rules you configure for a given URL pattern based on file extensions, file names, or any portion of a URL path on your website (e.g., *.jpg). You can configure multiple cache behaviors for your web distribution. Amazon CloudFront will match incoming viewer requests with your list of URL patterns, and if there is a match, the service will honor the cache behavior you configure for that URL pattern. Each cache behavior can include the following Amazon CloudFront configuration values: origin server name, viewer connection protocol, minimum expiration period, query string parameters, cookies, and trusted signers for private content.

Reference: <https://aws.amazon.com/Cloudfront/dynamic-content/>

NEW QUESTION 2

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 3

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

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

Answer: B

Explanation:

For the highest possible performance, all instances in a clustered database like this one should be in a single Availability Zone in a placement group, using EBS optimized instances, and using PIOPS SSD EBS Volumes. The particular Availability Zone the system is running in should not be important, as long as it is the same as the requesting resources.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>

NEW QUESTION 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, CloudTrail, IAM Credential Reports
- B. CloudTrail, IAM Credential Reports, AWS Config
- C. CloudTrail, AWS Config, IAM Credential Reports
- D. AWS Config, IAM Credential Reports, CloudTrail

Answer: C

Explanation:

You can use AWS CloudTrail 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/awsccloudtrail/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 CloudFront.
- C. Create read replicas for RDS since the load is mostly reads.
- D. Create a Multi-AZ RDS installs and route read traffic to standby

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 MariaDB Read Replicas](#).
Reference: <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZ.html>

NEW QUESTION 6

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 machine.
- D. Use these credits whenever you create AMIs on instances.
- E. When the CI instance receives commits, attach a new EBS volume to the CI machine.
- 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 times for a total of up to 50% savings.
Reference: <https://aws.amazon.com/ec2/spot/pricing/>

NEW QUESTION 7

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 for jobs to process.
- B. Use an SQS queue and set the reprocessing timeout to long enough for jobs to process.
- C. Use an SQS queue and set the visibility timeout to long enough for jobs to process.
- D. Use an SNS queue and set the reprocessing timeout to long enough for jobs to process.

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/AWSSimpleQueueService/latest/SQSDeveloperGuide/MessageLifecycle.html>

NEW QUESTION 8

You are designing a service that aggregates clickstream data in batch and delivers reports to subscribers via email only once per week. Data is extremely spiky, geographically distributed, high-scale, and unpredictable. How should you design this system?

- A. Use a large RedShift cluster to perform the analysis, and a fleet of Lambdas to perform record inserts into the RedShift table.
- B. Lambda will scale rapidly enough for the traffic spikes.
- C. Use a CloudFront distribution with access log delivery to S3. Clicks should be recorded as querystring GETs to the distribution.
- D. Reports are built and sent by periodically running EMR jobs over the access logs in S3.
- E. Use API Gateway invoking Lambdas which PutRecords into Kinesis, and EMR running Spark performing GetRecords on Kinesis to scale with spike.
- F. Spark on EMR outputs the analysis to S3, which are sent out via email.
- G. Use AWS Elasticsearch service and EC2 Auto Scaling group.
- H. The Autoscaling groups scale based on click throughput and stream into the Elasticsearch domain, which is also scalable.
- I. Use Kibana to generate reports periodically.

Answer: B

Explanation:

Because you only need to batch analyze, anything using streaming is a waste of money. CloudFront is a Gigabit-Scale HTTP(S) global request distribution service, so it can handle scale, geo-spread, spikes, and unpredictability. The Access Logs will contain the GET data and work just fine for batch analysis and email using EMR.

Can I use Amazon CloudFront if I expect usage peaks higher than 10 Gbps or 15,000 RPS? Yes. Complete our request for higher limits here, and we will add more capacity to your account within two business days.

Reference: <https://aws.amazon.com/Cloudfront/faqs/>

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 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 10

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

- A. Fn::Equals
- B. Fn::If
- C. Fn::Not
- D. Fn::Parse

Answer: D

Explanation:

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

NEW QUESTION 10

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

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

Answer: A

Explanation:

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

NEW QUESTION 12

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 `PutBucket` 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 CloudWatch 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 `PutBucket` 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 CloudWatch 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 14

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 software
- 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 15

Which EBS volume type is best for high performance NoSQL cluster deployments?

- A. io1
- B. gp1
- 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/EBSVolumeTypes.html>

NEW QUESTION 19

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 fleet 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 21

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-service.html>

NEW QUESTION 25

Which status represents a failure state in AWS CloudFormation?

- A. `UPDATE_COMPLETE_CLEANUP_IN_PROGRESS`
- B. `DELETE_COMPLETE_WITH_ARTIFACTS`
- C. `ROLLBACK_IN_PROGRESS`
- D. `ROLLBACK_FAILED`

Answer: C

Explanation:

ROLLBACK_IN_PROGRESS means an UpdateStack operation failed and the stack is in the process of trying to return to the valid, pre-update state.

UPDATE_COMPLETE_CLEANUP_IN_PROGRESS means an update was successful, and CloudFormation is deleting any replaced, no longer used resources.

ROLLBACK_FAILED is not a CloudFormation state (but UPDATE_ROLLBACK_FAILED is). DELETE_COMPLETE_WITH_ARTIFACTS does not exist at all.

Reference:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-updating-stacks.html>

NEW QUESTION 30

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 34

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 36

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

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

Answer: B

Explanation:

You can manage Amazon SQS messages with Amazon S3. This is especially useful for storing and retrieving messages with a message size of up to 2 GB. To manage Amazon SQS messages with Amazon S3, use the Amazon SQS Extended Client Library for Java.

Reference:

<http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/s3-messages.html>

NEW QUESTION 39

Which deployment method, when using AWS Auto Scaling Groups and Auto Scaling Launch Configurations, enables the shortest time to live for individual servers?

- 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 fleet

Answer: A

Explanation:

Note that the bootstrapping process can be slower if you have a complex application or multiple applications to install. Managing a fleet 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 42

Which major database needs a BYO license?

- A. PostgreSQL
- B. MariaDB
- 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 43

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

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

Answer: B

Explanation:

When a scale out event occurs, the Auto Scaling group launches the required number of EC2 instances, using its assigned launch configuration. These instances start in the Pending state. If you add a lifecycle hook to your Auto Scaling group, you can perform a custom action here. For more information, see Lifecycle Hooks.

Reference: <http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AutoScalingGroupLifecycle.html>

NEW QUESTION 44

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 subscription
- D. When a threshold is exceeded, email the manager.
- E. Scrape the billing page periodically and publish to SNS

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 48

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 hardware
- 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 53

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. Divide and round up for 5556.

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

NEW QUESTION 58

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 CloudTrail 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 time

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 60

For AWS CloudFormation, which is true?

- A. Custom resources using SNS have a default timeout of 3 minutes.
- B. Custom resources using SNS do not need a `ServiceToken` property.
- C. Custom resources using Lambda and `Code.ZipFile` allow inline nodejs resource composition.
- D. Custom resources using Lambda do not need a `ServiceToken` property

Answer: C

Explanation:

Code is a property of the AWS::Lambda::Function resource that enables to you specify the source code of an AWS Lambda (Lambda) function. You can point to a file in an Amazon Simple Storage Service (Amazon S3) bucket or specify your source code as inline text (for nodejs runtime environments only). Reference: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/template-custom-resources.html>

NEW QUESTION 63

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

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

Answer: B

Explanation:

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

NEW QUESTION 65

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/latest/dev/introduction.html>

NEW QUESTION 70

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 75

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 80

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

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

Answer: D

Explanation:

Making RDS accessible to the public internet in a public subnet poses a security risk, by making your database directly addressable and spamnable.

DB instances deployed within a VPC can be configured to be accessible from the Internet or from EC2 instances outside the VPC. If a VPC security group specifies a port access such as TCP port 22, you would not be able to access the DB instance because the firewall for the DB instance provides access only via the IP addresses specified by the DB security groups the instance is a member of and the port defined when the DB instance was created.

Reference: <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.RDSSecurityGroups.html>

NEW QUESTION 81

Which of these is not a reason a Multi-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.MultiAZ.html>

NEW QUESTION 86

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

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

Answer: C

Explanation:

Cost Allocation Tagging is a built-in feature of AWS, and when coupled with the Cost Explorer, provides a simple and robust way to track expenses.

You can also use tags to filter views in Cost Explorer. Note that before you can filter views by tags in Cost Explorer, you must have applied tags to your resources and activate them, as described in the following sections. For more information about Cost Explorer, see Analyzing Your Costs with Cost Explorer. Reference:

<http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/cost-alloc-tags.html>

NEW QUESTION 87

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 88

What is true of the way that encryption works with EBS?

- A. Snapshotting an encrypted volume makes an encrypted snapshot; restoring an encrypted snapshot creates an encrypted volume when specified / requested.
- B. Snapshotting an encrypted volume makes an encrypted snapshot when specified / requested; restoring an encrypted snapshot creates an encrypted volume when specified / requested.
- C. Snapshotting an encrypted volume makes an encrypted snapshot; restoring an encrypted snapshot always creates an encrypted volume.
- D. Snapshotting an encrypted volume makes an encrypted snapshot when specified / requested; restoring an encrypted snapshot always creates an encrypted volume.

Answer: C

Explanation:

Snapshots that are taken from encrypted volumes are automatically encrypted. Volumes that are created from encrypted snapshots are also automatically encrypted. Your encrypted volumes and any associated snapshots always remain protected. For more information, see Amazon EBS Encryption. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html>

NEW QUESTION 92

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 (IdP) -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/IAM/latest/UserGuide/id_roles_providers_oidc.html

NEW QUESTION 93

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

- A. DynamoDB table with 100x higher read than write throughput, with CloudFront caching.
- B. DynamoDB table with roughly equal read and write throughput, with CloudFront caching.
- C. DynamoDB table with 100x higher read than write throughput, with ElastiCache caching.
- D. DynamoDB table with roughly equal read and write throughput, with ElastiCache caching

Answer: D

Explanation:

Because the 100x read ratio is mostly driven by a small subset, with caching, only a roughly equal number of reads to writes will miss the cache, since the supermajority will hit the top 1% scores. Knowing we need to set the values roughly equal when using caching, we select AWS ElastiCache, because CloudFront cannot directly cache DynamoDB queries, and ElastiCache is an excellent in-memory cache for database queries, rather than a distributed proxy cache for content delivery.

One solution would be to cache these reads at the application layer. Caching is a technique that is used in many high-throughput applications, offloading read activity on hot items to the cache rather than to the database. Your application can cache the most popular items in memory, or use a product such as ElastiCache to do the same.

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

NEW QUESTION 96

You were just hired as a DevOps Engineer for a startup. Your startup uses AWS for 100% of their infrastructure. They currently have no automation at all for deployment, and they have had many failures while trying to deploy to production. The company has told you deployment process risk mitigation is the most important thing now, and you have a lot of budget for tools and AWS resources.

Their stack: 2-tier API

Data stored in DynamoDB or S3, depending on type Compute layer is EC2 in Auto Scaling Groups They use Route53 for DNS pointing to an ELB

An ELB balances load across the EC2 instances

The scaling group properly varies between 4 and 12 EC2 servers.

Which of the following approaches, given this company's stack and their priorities, best meets the company's needs?

- A. Model the stack in AWS Elastic Beanstalk as a single Application with multiple Environment
- B. Use Elastic Beanstalk's Rolling Deploy option to progressively roll out application code changes when promoting across environments.
- C. Model the stack in 3 CloudFormation templates: Data layer, compute layer, and networking layer
- D. Write stack deployment and integration testing automation following Blue-Green methodologies.

- E. Model the stack in AWS OpsWorks as a single Stack, with 1 compute layer and its associated EL
- F. Use Chef and App Deployments to automate Rolling Deployment.
- G. Model the stack in 1 CloudFormation template, to ensure consistency and dependency graph resolution
- H. Write deployment and integration testing automation following Rolling Deployment methodologies.

Answer: B

Explanation:

AWS recommends Blue-Green for zero-downtime deploys. Since you use DynamoDB, and neither AWS OpsWorks nor AWS Elastic Beanstalk directly supports DynamoDB, the option selecting CloudFormation and Blue-Green is correct.

You use various strategies to migrate the traffic from your current application stack (blue) to a new version of the application (green). This is a popular technique for deploying applications with zero downtime. The deployment services like AWS Elastic Beanstalk, AWS CloudFormation, or AWS OpsWorks are particularly useful as they provide a simple way to clone your running application stack. You can set up a new version of your application (green) by simply cloning current version of the application (blue). Reference: <https://d0.awsstatic.com/whitepapers/overview-of-deployment-options-on-aws.pdf>

NEW QUESTION 100

Your system uses a multi-master, multi-region DynamoDB configuration spanning two regions to achieve high availability. 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 check 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.CrossRegionRepl.html>

NEW QUESTION 103

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 logical layers 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 entries 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 104

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 capacity
- 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, to justify adding capacity instead of wasting money.

Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/policy_creating.html

NEW QUESTION 108

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 112

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 117

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 119

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

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

- * AWS-Certified-DevOps-Engineer-Professional Questions and Answers Updated Frequently
- * AWS-Certified-DevOps-Engineer-Professional Practice Questions Verified by Expert Senior Certified Staff
- * AWS-Certified-DevOps-Engineer-Professional Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * AWS-Certified-DevOps-Engineer-Professional Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

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