

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 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 3

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 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 3rd

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

What is the scope of an EC2 security group?

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

Answer: C

Explanation:

A security group is tied to a region and can be assigned only to instances in the same region. You can't enable an instance to communicate with an instance outside its region using security group rules. Traffic from an instance in another region is seen as WAN bandwidth.

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

NEW QUESTION 6

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 software
- 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 software
- 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 software
- 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 7

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 8

When thinking of AWS Elastic Beanstalk'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/Welcome.html>

NEW QUESTION 9

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

- A. Subscribe your queue to an SNS topic instead.
- B. Use as long of a poll as possible, instead of short polls.
- C. Alter your visibility timeout to be shorter.
- D. Use `qsqs` on your EC2 instance

Answer: B

Explanation:

One benefit of long polling with Amazon SQS is the reduction of the number of empty responses, when there are no messages available to return, in reply to a ReceiveMessage request sent to an Amazon SQS queue. Long polling allows the Amazon SQS service to wait until a message is available in the queue before sending a response.

Reference:

<http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-long-polling.html>

NEW QUESTION 10

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 10

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 table.

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 "global" 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 12

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

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

Answer: C

Explanation:

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

Reference:

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

NEW QUESTION 17

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.html

NEW QUESTION 18

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 22

For AWS CloudFormation, which stack state refuses UpdateStack calls?

- A. `UPDATE_ROLLBACK_FAILED`
- B. `UPDATE_ROLLBACK_COMPLETE`
- C. `UPDATE_CONIplete`
- D. `CREATE_COMPLETE`

Answer: A

Explanation:

When a stack is in the UPDATE_ROLLBACK_FAILED 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_FAILED 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-continueupdateandrollback.html>

NEW QUESTION 23

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 28

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 CloudTrail 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 activity. 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-file-validation-intro.html>

NEW QUESTION 29

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 30

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 31

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 Microsoft Active Directory, without caching any information in the cloud. AD Connector comes in 2 sizes; small and large. A small AD Connector is designed for smaller organizations of up to 500 users. A large AD Connector is designed for larger organizations of up to 5,000 users.

Reference: <https://aws.amazon.com/directoryservice/details/>

NEW QUESTION 35

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 36

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 40

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 CloudTrail Trail
- 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 table
- G. Generate reports based on the contents of this table.

Answer: A

Explanation:

This is the ideal use case for AWS CloudTrail.

CloudTrail provides visibility into user activity by recording API calls made on your account. CloudTrail 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. CloudTrail makes it easier to ensure compliance with internal policies and regulatory standards. Reference: <https://aws.amazon.com/Cloudtrail/faqs/>

NEW QUESTION 44

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 bplex 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 47

If I want CloudFormation 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 CloudFormation stack and display those state changes in the UI for the system.
- B. Use a long-poll on the `ListStacks` API call for your CloudFormation 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 CloudFormation stack to publish events into.
- D. Subscribe your continuous delivery system to an SQS queue that you also tell your CloudFormation 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/AWSCloudFormation/latest/UserGuide/using-cfn-updating-stacks-monitor-stack.html>

NEW QUESTION 52

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 OS

Answer: C

Explanation:

This is the definition of Unavailable from the EC2 and EBS SLA.

"Unavailable" and "Unavailability" 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/sla/>

NEW QUESTION 56

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 58

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

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

Answer: D

Explanation:

Making RDS accessible to the public internet in a public subnet poses a security risk, by making your database directly addressable and spammable. DB instances deployed within a VPC can be configured to be accessible from the Internet or from EC2 instances outside the VPC. If a VPC security group specifies a port access such as TCP port 22, you would not be able to access the DB instance because the firewall for the DB instance provides access only via the IP addresses specified by the DB security groups the instance is a member of and the port defined when the DB instance was created.
 Reference: <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.RDSSecurityGroups.html>

NEW QUESTION 62

You need your API backed by DynamoDB to stay online during a total regional AWS failure. You can tolerate a couple minutes of lag or slowness during a large failure event, but the system should recover with normal operation after those few minutes. What is a good approach?

- A. Set up DynamoDB cross-region replication in a master-standby configuration, with a single standby in another region
- B. Create an Auto Scaling Group behind an ELB in each of the two regions DynamoDB is running in
- C. Add a Route53 Latency DNS Record with DNS Failover, using the ELBs in the two regions as the resource records.
- D. Set up a DynamoDB Multi-Region table
- E. Create an Auto Scaling Group behind an ELB in each of the two regions DynamoDB is running in
- F. Add a Route53 Latency DNS Record with DNS Failover, using the ELBs in the two regions as the resource records.
- G. Set up a DynamoDB Multi-Region table
- H. Create a cross-region ELB pointing to a cross-region Auto Scaling Group, and direct a Route53 Latency DNS Record with DNS Failover to the cross-region ELB.
- I. Set up DynamoDB cross-region replication in a master-standby configuration, with a single standby in another region
- J. Create a cross-region ELB pointing to a cross-region Auto Scaling Group, and direct a Route53 Latency DNS Record with DNS Failover to the cross-region ELB.

Answer: A

Explanation:

There is no such thing as a cross-region ELB, nor such thing as a cross-region Auto Scaling Group, nor such thing as a DynamoDB Multi-Region Table. The only option that makes sense is the cross-regional replication version with two ELBs and ASGs with Route53 Failover and Latency DNS.
 Reference: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.CrossRegionRepl.html>

NEW QUESTION 66

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 it
- 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 71

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

- A. AWS Config
- B. Amazon CloudWatch Metrics
- 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 72

What is a circular dependency in AWS CloudFormation?

- 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 Template.

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 CloudFormation 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.html#troubleshooting-g-errors-dependence-error>

NEW QUESTION 75

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 CloudFormation. 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/AWSCloudFormation/latest/UserGuide/aws-resource-emr-cluster.html>

NEW QUESTION 76

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 CloudFormation.
- 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 parity.
- 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 virtual 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 clustering.
- G. This will make sure that the AMIs and machine sizes are the same across both environments.

Answer: A

Explanation:

Only CloudFormation's JSON Templates allow declarative version control of repeatably deployable models of entire AWS clouds.

Reference: <https://blogs.aws.amazon.com/application-management/blog/category/Best+practices>

NEW QUESTION 79

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 83

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 88

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 91

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 94

.....

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 First Try
- * AWS-Certified-DevOps-Engineer-Professional Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

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