



Amazon

Exam Questions AWS-Certified-DevOps-Engineer-Professional

Amazon AWS Certified DevOps Engineer Professional

NEW QUESTION 1

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

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

Answer: B

Explanation:

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

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

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 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 4

When thinking of AWS Elastic Beanstalk, the 'Swap Environment URLs' feature most directly aids in what?

- A. Immutable Rolling Deployments
- B. Mutable Rolling Deployments
- C. Canary Deployments
- D. Blue-Green Deployments

Answer: D

Explanation:

Simply upload the new version of your application and let your deployment service (AWS Elastic Beanstalk, AWS CloudFormation, or AWS OpsWorks) deploy a new version (green). To cut over to the new version, you simply replace the ELB URLs in your DNS records. Elastic Beanstalk has a Swap Environment URLs feature to facilitate a simpler cutover process.

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

NEW QUESTION 5

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 6

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 7

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 8

To monitor API calls against our AWS account by different users and entities, we can use to create a history of calls in bulk for later review, and use for reacting to AWS API calls in real-time.

- A. AWS Config; AWS Inspector
- B. AWS CloudTrail; AWS Config
- C. AWS CloudTrail; CloudWatch Events
- D. AWS Config; AWS Lambda

Answer: C

Explanation:

CloudTrail is a batch API call collection service, CloudWatch Events enables real-time monitoring of calls through the Rules object interface.

Reference: <https://aws.amazon.com/whitepapers/security-at-scale-governance-in-aws/>

NEW QUESTION 9

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.MultiAZ.html>

NEW QUESTION 10

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

- A. Copy all log files into AWS S3 using a cron job on each instance
- B. Use an S3 Notification Configuration on the `PutBucket` event and publish events to AWS Lambda
- C. Use the Lambda to analyze logs as soon as they come in and flag issues.
- D. Begin using CloudWatch Logs on every service
- E. Stream all Log Groups into S3 object
- F. Use AWS EMR cluster jobs to perform ad-hoc MapReduce analysis and write new queries when needed.
- G. Copy all log files into AWS S3 using a cron job on each instance
- H. Use an S3 Notification Configuration on the `PutBucket` event and publish events to AWS Kinesis
- I. Use Apache Spark on AWS EMR to perform at-scale stream processing queries on the log chunks and flag issues.
- J. Begin using CloudWatch Logs on every service
- K. Stream all Log Groups into an AWS Elasticsearch Service Domain running Kibana 4 and perform log analysis on a search cluster.

Answer: D

Explanation:

The Elasticsearch and Kibana 4 combination is called the ELK Stack, and is designed specifically for real-time, ad-hoc log analysis and aggregation. All other answers introduce extra delay or require pre-defined queries.

Amazon Elasticsearch Service is a managed service that makes it easy to deploy, operate, and scale Elasticsearch in the AWS Cloud. Elasticsearch is a popular open-source search and analytics engine for use cases such as log analytics, real-time application monitoring, and click stream analytics. Reference: <https://aws.amazon.com/elasticsearch-service/>

NEW QUESTION 10

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 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 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 as a "Full" level utilization purchases on EC2.

Reference: https://d0.awsstatic.com/whitepapers/Cost_Optimization_with_AWS.pdf

NEW QUESTION 19

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 24

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 29

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 32

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 33

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 34

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 36

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 40

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 45

You are building a deployment system on AWS. You will deploy new code by bootstrapping instances in a private subnet in a VPC at runtime using UserData scripts pointing to an S3 zip file object, where your code is stored. An ELB in a public subnet has network interfaces and connects to the instances. Requests from users of the system are routed to the ELB via a Route53 A Record Alias. You do not use any VPC endpoints. Which is a risk of using this approach?

- A. Route53 Alias records do not always update dynamically with ELB network changes after deploys.
- B. If the NAT routing for the private subnet fails, deployments fail.
- C. Kernel changes to the base AMI may render the code inoperable.
- D. The instances cannot be in a private subnet if the ELB is in a public one

Answer: B

Explanation:

Since you are not using VPC endpoints, outbound requests for the code sitting in S3 are routed through the NAT for the VPC's private subnets. If this networking fails, runtime bootstrapping through code

download will fail due to network unavailability and lack of access to the Internet, and thus Amazon S3. Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_NAT_Instance.html

NEW QUESTION 48

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 51

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 52

For AWS CloudFormation, which stack state refuses UpdateStack calls?

- A. `UPDATE_ROLLBACK_FAILED`
- B. `UPDATE_ROLLBACK_COMPLETE`
- C. `UPDATE_COMPLETE`
- 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-continue-updaterollback.html>

NEW QUESTION 56

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 57

Which of these is not a Pseudo Parameter in AWS CloudFormation?

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

Answer: C

Explanation:

This is the complete list of Pseudo Parameters: `AWS::AccountId`, `AWS::NotificationARNs`, `AWS::NoValue`, `AWS::Region`, `AWS::StackId`, `AWS::StackName`

Reference:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/pseudo-parameter-reference.html>

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 you to 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 61

You need to deploy an AWS stack in a repeatable manner across multiple environments. You have selected CloudFormation as the right tool to accomplish this, but have found that there is a resource type you need to create and model, but is unsupported by CloudFormation. How should you overcome this challenge?

- A. Use a CloudFormation Custom Resource Template by selecting an API call to proxy for create, update, and delete action
- B. CloudFormation will use the AWS SDK, CLI, or API method of your choosing as the state transition function for the resource type you are modeling.
- C. Submit a ticket to the AWS Forum
- D. AWS extends CloudFormation Resource Types by releasing tooling to the AWS Labs organization on GitHub
- E. Their response time is usually 1 day, and they complete requests within a week or two.
- F. Instead of depending on CloudFormation, use Chef, Puppet, or Ansible to author Heat templates, which are declarative stack resource definitions that operate over the OpenStack hypervisor and cloud environment.
- G. Create a CloudFormation Custom Resource Type by implementing create, update, and delete functionality, either by subscribing a Custom Resource Provider to an SNS topic, or by implementing the logic in AWS Lambda.

Answer: D

Explanation:

Custom resources provide a way for you to write custom provisioning logic in AWS CloudFormation template and have AWS CloudFormation run it during a stack operation, such as when you create, update or delete a stack. For more information, see Custom Resources.

Reference:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/template-custom-resources.html>

NEW QUESTION 66

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 71

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 75

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 78

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 Configuration
- B. Create AMIs with all code pre-installed
- C. Assign the new AMI to the second Auto Scaling Launch Configuration
- 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 needed
- 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-installed
- H. Assign the new AMI to the Auto Scaling Launch Configuration, to replace the old one
- 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 code
- 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 Beanstalk
- 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 time
- M. Re-deploy the old code bundle to rollback if needed.

Answer: A

Explanation:

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

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

NEW QUESTION 83

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 84

You need to create an audit log of all changes to customer banking data. You use DynamoDB to store this customer banking data. It's important not to lose any information due to server failures. What is an elegant way to accomplish this?

- A. Use a DynamoDB StreamSpecification and stream all changes to AWS Lambda
- B. Log the changes to AWS CloudWatch Logs, removing sensitive information before logging.
- C. Before writing to DynamoDB, do a pre-write acknowledgment to disk on the application server, removing sensitive information before logging
- D. Periodically rotate these log files into S3.
- E. Use a DynamoDB StreamSpecification and periodically flush to an EC2 instance store, removing sensitive information before putting the object
- F. Periodically flush these batches to S3.
- G. Before writing to DynamoDB, do a pre-write acknowledgment to disk on the application server, removing sensitive information before logging
- H. Periodically pipe these files into CloudWatch Logs.

Answer: A

Explanation:

All suggested periodic options are sensitive to server failure during or between periodic flushes. Streaming to Lambda and then logging to CloudWatch Logs will make the system resilient to instance and Availability Zone failures.

Reference: <http://docs.aws.amazon.com/lambda/latest/dg/with-ddb.html>

NEW QUESTION 85

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 88

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 93

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 98

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 layer
- D. When the instances are booted and the HTTP APIs are up, submit processing requests to individual instances.
- E. Make an S3 notification configuration which publishes to AWS Lambda on the manifest bucket
- F. Make the Lambda create a CloudFormation Stack which contains the logic to construct an autoscaling worker tier of EC2 G2 instances with the ANN code on each instance
- G. Create an SQS queue of the images in the manifest
- H. Tear the stack down when the queue is empty.
- I. Deploy your ANN code to AWS Lambda as a bundled binary for the C++ extension
- J. Make an S3 notification configuration on the manifest, which publishes to another AWS Lambda running controller code
- K. This controller code publishes all the images in the manifest to AWS Kinesis
- L. Your ANN code Lambda Function uses the Kinesis as an Event Source
- M. The system automatically scales when the stream contains image events.
- N. Create an Auto Scaling, Load Balanced Elastic Beanstalk worker tier Application and Environment
- O. Deploy the ANN code to G2 instances in this tier
- P. Set the desired capacity to 1. Make the code periodically check S3 for new manifest
- Q. When a new manifest is detected, push all of the images in the manifest into the SQS queue associated with the Elastic Beanstalk worker tier.

Answer: B

Explanation:

The Elastic Beanstalk option is incorrect because it requires a constantly-polling instance, which may break and costs money.

The Lambda fleet option is incorrect because AWS Lambda does not support GPU usage.

The OpsWorks stack option both requires a constantly-polling instance, and also requires complex timing and capacity planning logic.

The CloudFormation option requires no polling, has no always-on instances, and allows arbitrarily fast processing by simply setting the instance count as high as needed.

Reference: <http://docs.aws.amazon.com/lambda/latest/dg/current-supported-versions.html>

NEW QUESTION 103

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

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

Answer: C

Explanation:

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

Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html#concepts-regions-availability-zones>

NEW QUESTION 108

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 user.

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 109

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 114

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 instances.

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 Environments.
- 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 ELB.
- 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->

NEW QUESTION 117

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 118

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 122

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 125

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 do not have noisy neighbors.
- D. Turn off SYN/ACK on your TCP stack or begin using UDP for higher throughput.

Answer: A

Explanation:

For instances that are 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.html#jumbo_frame_instances

NEW QUESTION 127

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 account

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 132

.....

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

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

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

Answer: B

Explanation:

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

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

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 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 4

When thinking of AWS Elastic Beanstalk, the 'Swap Environment URLs' feature most directly aids in what?

- A. Immutable Rolling Deployments
- B. Mutable Rolling Deployments
- C. Canary Deployments
- D. Blue-Green Deployments

Answer: D

Explanation:

Simply upload the new version of your application and let your deployment service (AWS Elastic Beanstalk, AWS CloudFormation, or AWS OpsWorks) deploy a new version (green). To cut over to the new version, you simply replace the ELB URLs in your DNS records. Elastic Beanstalk has a Swap Environment URLs feature to facilitate a simpler cutover process.

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

NEW QUESTION 5

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 6

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 7

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 8

To monitor API calls against our AWS account by different users and entities, we can use to create a history of calls in bulk for later review, and use for reacting to AWS API calls in real-time.

- A. AWS Config; AWS Inspector
- B. AWS CloudTrail; AWS Config
- C. AWS CloudTrail; CloudWatch Events
- D. AWS Config; AWS Lambda

Answer: C

Explanation:

CloudTrail is a batch API call collection service, CloudWatch Events enables real-time monitoring of calls through the Rules object interface.

Reference: <https://aws.amazon.com/whitepapers/security-at-scale-governance-in-aws/>

NEW QUESTION 9

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.MultiAZ.html>

NEW QUESTION 10

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

- A. Copy all log files into AWS S3 using a cron job on each instance
- B. Use an S3 Notification Configuration on the `PutBucket` event and publish events to AWS Lambda
- C. Use the Lambda to analyze logs as soon as they come in and flag issues.
- D. Begin using CloudWatch Logs on every service
- E. Stream all Log Groups into S3 object
- F. Use AWS EMR cluster jobs to perform ad-hoc MapReduce analysis and write new queries when needed.
- G. Copy all log files into AWS S3 using a cron job on each instance
- H. Use an S3 Notification Configuration on the `PutBucket` event and publish events to AWS Kinesis
- I. Use Apache Spark on AWS EMR to perform at-scale stream processing queries on the log chunks and flag issues.
- J. Begin using CloudWatch Logs on every service
- K. Stream all Log Groups into an AWS Elasticsearch Service Domain running Kibana 4 and perform log analysis on a search cluster.

Answer: D

Explanation:

The Elasticsearch and Kibana 4 combination is called the ELK Stack, and is designed specifically for real-time, ad-hoc log analysis and aggregation. All other answers introduce extra delay or require pre-defined queries.

Amazon Elasticsearch Service is a managed service that makes it easy to deploy, operate, and scale Elasticsearch in the AWS Cloud. Elasticsearch is a popular open-source search and analytics engine for use cases such as log analytics, real-time application monitoring, and click stream analytics. Reference: <https://aws.amazon.com/elasticsearch-service/>

NEW QUESTION 10

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 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 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 as a "Full" level utilization purchases on EC2.

Reference: https://d0.awsstatic.com/whitepapers/Cost_Optimization_with_AWS.pdf

NEW QUESTION 19

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 24

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 29

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 32

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 33

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 34

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 36

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 40

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 45

You are building a deployment system on AWS. You will deploy new code by bootstrapping instances in a private subnet in a VPC at runtime using UserData scripts pointing to an S3 zip file object, where your code is stored. An ELB in a public subnet has network interfaces and connects to the instances. Requests from users of the system are routed to the ELB via a Route53 A Record Alias. You do not use any VPC endpoints. Which is a risk of using this approach?

- A. Route53 Alias records do not always update dynamically with ELB network changes after deploys.
- B. If the NAT routing for the private subnet fails, deployments fail.
- C. Kernel changes to the base AMI may render the code inoperable.
- D. The instances cannot be in a private subnet if the ELB is in a public one

Answer: B

Explanation:

Since you are not using VPC endpoints, outbound requests for the code sitting in S3 are routed through the NAT for the VPC's private subnets. If this networking fails, runtime bootstrapping through code

download will fail due to network unavailability and lack of access to the Internet, and thus Amazon S3. Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_NAT_Instance.html

NEW QUESTION 48

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 51

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 52

For AWS CloudFormation, which stack state refuses UpdateStack calls?

- A. `UPDATE_ROLLBACK_FAILED`
- B. `UPDATE_ROLLBACK_COMPLETE`
- C. `UPDATE_COMPLETE`
- 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-continue-updaterollback.html>

NEW QUESTION 56

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 57

Which of these is not a Pseudo Parameter in AWS CloudFormation?

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

Answer: C

Explanation:

This is the complete list of Pseudo Parameters: `AWS::AccountId`, `AWS::NotificationARNs`, `AWS::NoValue`, `AWS::Region`, `AWS::StackId`, `AWS::StackName`

Reference:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/pseudo-parameter-reference.html>

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 you to 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 61

You need to deploy an AWS stack in a repeatable manner across multiple environments. You have selected CloudFormation as the right tool to accomplish this, but have found that there is a resource type you need to create and model, but is unsupported by CloudFormation. How should you overcome this challenge?

- A. Use a CloudFormation Custom Resource Template by selecting an API call to proxy for create, update, and delete action
- B. CloudFormation will use the AWS SDK, CLI, or API method of your choosing as the state transition function for the resource type you are modeling.
- C. Submit a ticket to the AWS Forum
- D. AWS extends CloudFormation Resource Types by releasing tooling to the AWS Labs organization on GitHub
- E. Their response time is usually 1 day, and they complete requests within a week or two.
- F. Instead of depending on CloudFormation, use Chef, Puppet, or Ansible to author Heat templates, which are declarative stack resource definitions that operate over the OpenStack hypervisor and cloud environment.
- G. Create a CloudFormation Custom Resource Type by implementing create, update, and delete functionality, either by subscribing a Custom Resource Provider to an SNS topic, or by implementing the logic in AWS Lambda.

Answer: D

Explanation:

Custom resources provide a way for you to write custom provisioning logic in AWS CloudFormation template and have AWS CloudFormation run it during a stack operation, such as when you create, update or delete a stack. For more information, see Custom Resources.

Reference:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/template-custom-resources.html>

NEW QUESTION 66

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 71

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 75

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 78

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 Configuration
- B. Create AMIs with all code pre-installed
- C. Assign the new AMI to the second Auto Scaling Launch Configuration
- 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 needed
- 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-installed
- H. Assign the new AMI to the Auto Scaling Launch Configuration, to replace the old one
- 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 code
- 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 Beanstalk
- 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 time
- M. Re-deploy the old code bundle to rollback if needed.

Answer: A

Explanation:

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

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

NEW QUESTION 83

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 84

You need to create an audit log of all changes to customer banking data. You use DynamoDB to store this customer banking data. It's important not to lose any information due to server failures. What is an elegant way to accomplish this?

- A. Use a DynamoDB StreamSpecification and stream all changes to AWS Lambda
- B. Log the changes to AWS CloudWatch Logs, removing sensitive information before logging.
- C. Before writing to DynamoDB, do a pre-write acknowledgment to disk on the application server, removing sensitive information before logging
- D. Periodically rotate these log files into S3.
- E. Use a DynamoDB StreamSpecification and periodically flush to an EC2 instance store, removing sensitive information before putting the object
- F. Periodically flush these batches to S3.
- G. Before writing to DynamoDB, do a pre-write acknowledgment to disk on the application server, removing sensitive information before logging
- H. Periodically pipe these files into CloudWatch Logs.

Answer: A

Explanation:

All suggested periodic options are sensitive to server failure during or between periodic flushes. Streaming to Lambda and then logging to CloudWatch Logs will make the system resilient to instance and Availability Zone failures.

Reference: <http://docs.aws.amazon.com/lambda/latest/dg/with-ddb.html>

NEW QUESTION 85

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 88

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 93

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 98

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 layer
- D. When the instances are booted and the HTTP APIs are up, submit processing requests to individual instances.
- E. Make an S3 notification configuration which publishes to AWS Lambda on the manifest bucket
- F. Make the Lambda create a CloudFormation Stack which contains the logic to construct an autoscaling worker tier of EC2 G2 instances with the ANN code on each instance
- G. Create an SQS queue of the images in the manifest
- H. Tear the stack down when the queue is empty.
- I. Deploy your ANN code to AWS Lambda as a bundled binary for the C++ extension
- J. Make an S3 notification configuration on the manifest, which publishes to another AWS Lambda running controller code
- K. This controller code publishes all the images in the manifest to AWS Kinesis
- L. Your ANN code Lambda Function uses the Kinesis as an Event Source
- M. The system automatically scales when the stream contains image events.
- N. Create an Auto Scaling, Load Balanced Elastic Beanstalk worker tier Application and Environment
- O. Deploy the ANN code to G2 instances in this tier
- P. Set the desired capacity to 1. Make the code periodically check S3 for new manifest
- Q. When a new manifest is detected, push all of the images in the manifest into the SQS queue associated with the Elastic Beanstalk worker tier.

Answer: B

Explanation:

The Elastic Beanstalk option is incorrect because it requires a constantly-polling instance, which may break and costs money.

The Lambda fleet option is incorrect because AWS Lambda does not support GPU usage.

The OpsWorks stack option both requires a constantly-polling instance, and also requires complex timing and capacity planning logic.

The CloudFormation option requires no polling, has no always-on instances, and allows arbitrarily fast processing by simply setting the instance count as high as needed.

Reference: <http://docs.aws.amazon.com/lambda/latest/dg/current-supported-versions.html>

NEW QUESTION 103

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

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

Answer: C

Explanation:

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

Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html#concepts-regions-availability-zones>

NEW QUESTION 108

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 user.

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 109

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 114

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 instances.

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 Environments.
- 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 ELB.
- 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->

NEW QUESTION 117

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 118

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 122

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 125

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 do not have noisy neighbors.
- D. Turn off SYN/ACK on your TCP stack or begin using UDP for higher throughput.

Answer: A

Explanation:

For instances that are 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.html#jumbo_frame_instances

NEW QUESTION 127

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 account

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 132

.....

Relate Links

100% Pass Your AWS-Certified-DevOps-Engineer-Professional Exam with ExamBible Prep Materials

<https://www.exambible.com/AWS-Certified-DevOps-Engineer-Professional-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>