



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

What is server immutability?

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

Answer: A

Explanation:

disposable upgrades offer a simpler way to know if your application has unknown dependencies. The underlying EC2 instance usage is considered temporary or ephemeral in nature for the period of deployment until the current release is active. During the new release, a new set of EC2 instances are rolled out by terminating older instances. This type of upgrade technique is more common in an immutable infrastructure.

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

NEW QUESTION 3

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 4

Fill the blanks: helps us track AWS API calls and transitions, helps to understand what resources we have now, and allows auditing credentials and logins.

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

Answer: C

Explanation:

You can use AWS CloudTrail to get a history of AWS API calls and related events for your account. This includes calls made by using the AWS Management Console, AWS SDKs, command line tools, and higher-level AWS services.

Reference: <http://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-user-guide.html>

NEW QUESTION 5

You are creating an application which stores extremely sensitive financial information. All information in the system must be encrypted at rest and in transit. Which of these is a violation of this policy?

- A. ELB SSL termination.
- B. ELB Using Proxy Protocol v1.
- C. CloudFront Viewer Protocol Policy set to HTTPS redirection.
- D. Telling S3 to use AES256 on the server-side

Answer: A

Explanation:

Terminating SSL terminates the security of a connection over HTTP, removing the S for "Secure" in HTTPS. This violates the "encryption in transit" requirement in the scenario.

Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/elb-listener-config.html>

NEW QUESTION 6

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 7

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 8

You need to process long-running jobs once and only once. How might you do this?

- A. Use an SNS queue and set the visibility timeout to long enough 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 9

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 10

Your system automatically provisions EIPs to EC2 instances in a VPC on boot. The system provisions the whole VPC and stack at once. You have two of them per VPC. On your new AWS account, your attempt to create a Development environment failed, after successfully creating Staging and Production environments in the same region. What happened?

- A. You didn't choose the Development version of the AMI you are using.
- B. You didn't set the Development flag to true when deploying EC2 instances.
- C. You hit the soft limit of 5 EIPs per region and requested a 6th.
- D. You hit the soft limit of 2 VPCs per region and requested a 3rd.

Answer: C

Explanation:

There is a soft limit of 5 EIPs per Region for VPC on new accounts. The third environment could not allocate the 6th EIP.

Reference: http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#limits_vpc

NEW QUESTION 10

How does Amazon RDS multi Availability Zone model work?

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

Answer: A

Explanation:

In a 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 12

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

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

Answer: A

Explanation:

When Auto Scaling responds to a scale in event, it terminates one or more instances. These instances are detached from the Auto Scaling group and enter the Terminating state.

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

NEW QUESTION 13

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 16

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

- A. io1
- B. gp1
- C. standard
- D. gp2

Answer: A

Explanation:

io1 volumes, or Provisioned IOPS (PIOPS) SSDs, are best for: Critical business applications that require sustained IOPS performance, or more than 10,000 IOPS or 160 MiB/s of throughput per volume, like large database workloads, such as MongoDB.

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

NEW QUESTION 18

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

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

Answer: B

Explanation:

the bootstrapping process can be slower if you have a complex application or multiple applications to install. Managing a fleet of applications with several build tools and dependencies can be a challenging task during rollouts. Furthermore, your deployment service should be designed to do faster rollouts to take advantage of Auto Scaling.

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

NEW QUESTION 20

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

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

Answer: A

Explanation:

Amazon Elasticsearch Service (Amazon ES) is a managed service that makes it easy to deploy, operate, and scale Elasticsearch clusters in the AWS cloud. Elasticsearch is a popular open-source search and analytics engine for use cases such as log analytics, real-time application monitoring, and click stream analytics.

Reference:

<http://docs.aws.amazon.com/elasticsearch-service/latest/developerguide/what-is-amazon-elasticsearch-service.html>

NEW QUESTION 25

What is the scope of an EC2 EIP?

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

Answer: C

Explanation:

An Elastic IP address is tied to a region and can be associated only with an instance in the same region. Reference:

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

NEW QUESTION 29

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

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

Answer: B

Explanation:

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

Reference:

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

NEW QUESTION 31

There are a number of ways to purchase compute capacity on AWS. Which orders the price per compute or memory unit from LOW to HIGH (cheapest to most expensive), on average?

(A) On-Demand (B) Spot (C) Reserved

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

Answer:

C

Explanation:

Spot instances are usually many, many times cheaper than on-demand prices. Reserved instances, depending on their term and utilization, can yield approximately 33% to 66% cost savings. On-Demand prices are the baseline price and are the most expensive way to purchase EC2 compute time. Reference: https://d0.awsstatic.com/whitepapers/Cost_Optimization_with_AWS.pdf

NEW QUESTION 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. IV|u|ti-AZ deployments do not provide sufficient availability when a business loses USD 360,000 per hour of unavailability. As RDS does not natively support multi-region, and ACID does not perform well/at all over large distances between regions, only the DynamoDB answer works. Reference: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.CrossRegionRepl.html>

NEW QUESTION 34

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 37

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

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

Answer: C

Explanation:

Linux stacks can install custom cookbooks from any of the following repository types: HTTP or Amazon S3 archives. They can be either public or private, but Amazon S3 is typically the preferred option for a private archive. Git and Subversion repositories provide source control and the ability to have multiple versions. Reference: <http://docs.aws.amazon.com/opsworks/latest/userguide/workingcookbook-installingcustom-enable.html>

NEW QUESTION 40

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 connectMty 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 on

Answer: B

Explanation:

Since you are not using VPC endpoints, outbound requests for the code sitting in S3 are routed though 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 41

Which major database needs a BYO license?

- A. PostgreSQL
- B. MariaDB
- C. MySQL
- D. Oracle

Answer: D

Explanation:

Oracle is not open source, and requires a bring your own license model.

Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Oracle.htm

NEW QUESTION 46

Why are more frequent snapshots or EBS Volumes faster?

- A. Blocks in EBS Volumes are allocated lazily, since while logically separated from other EBS Volumes, Volumes often share the same physical hardware
- B. Snapshotting the first time forces full block range allocation, so the second snapshot doesn't need to perform the allocation phase and is faster.
- C. The snapshots are incremental so that only the blocks on the device that have changed after your last snapshot are saved in the new snapshot.
- D. AWS provisions more disk throughput for burst capacity during snapshots if the drive has been pre-warmed by snapshotting and reading all blocks.
- E. The drive is pre-warmed, so block access is more rapid for volumes when every block on the device has already been read at least one time.

Answer: B

Explanation:

After writing data to an EBS volume, you can periodically create a snapshot of the volume to use as a baseline for new volumes or for data backup. If you make periodic snapshots of a volume, the snapshots are incremental so that only the blocks on the device that have changed after your last snapshot are saved in the new snapshot. Even though snapshots are saved incrementally, the snapshot deletion process is designed so that you need to retain only the most recent snapshot in order to restore the volume.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-creating-snapshot.html>

NEW QUESTION 47

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

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

Answer: A

Explanation:

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

Reference:

<http://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html>

NEW QUESTION 49

What is the scope of an EBS volume?

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

Answer: D

Explanation:

An Amazon EBS volume is tied to its Availability Zone and can be attached only to instances in the same Availability Zone.

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

NEW QUESTION 53

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

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

Answer: A

Explanation:

This is the complete list of Intrinsic Functions...: Fn::Base64, Fn::And, Fn::Equals, Fn::If, Fn::Not, Fn::Or, Fn::FindInMap, Fn::GetAtt, Fn::GetAZs, Fn::Join, Fn::Select, Ref

Reference:

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

NEW QUESTION 55

For AWS CloudFormation, which is true?

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

Answer: C

Explanation:

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

NEW QUESTION 57

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 59

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

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

Answer: B

Explanation:

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

NEW QUESTION 63

What is the scope of AWS IAM?

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

Answer: A

Explanation:

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

NEW QUESTION 68

Your CTO has asked you to make sure that you know what all users of your AWS account are doing to change resources at all times. She wants a report of who is doing what over time, reported to her once per week, for as broad a resource type group as possible. How should you do this?

- A. Create a global AWS CloudTrail Trail
- B. Configure a script to aggregate the log data delivered to S3 once per week and deliver this to the CTO.
- C. Use CloudWatch Events Rules with an SNS topic subscribed to all AWS API call
- D. Subscribe the CTO to an email type delivery on this SNS Topic.
- E. Use AWS IAM credential reports to deliver a CSV of all uses of IAM User Tokens over time to the CTO.
- F. Use AWS Config with an SNS subscription on a Lambda, and insert these changes over time into a DynamoDB table
- G. Generate reports based on the contents of this table.

Answer:

A

Explanation:

This is the ideal use case for AWS CloudTrail.

CloudTrail provides visibility into user activity by recording API calls made on your account. CloudTrail records important information about each API call, including the name of the API, the identity of the caller, the time of the API call, the request parameters, and the response elements returned by the AWS service. This information helps you to track changes made to your AWS resources and to troubleshoot operational issues. CloudTrail makes it easier to ensure compliance with internal policies and regulatory standards. Reference: <https://aws.amazon.com/Cloudtrail/faqs/>

NEW QUESTION 73

What is the order of most-to-least rapidly-scaling (fastest to scale first)?

(A) EC2 + ELB + Auto Scaling (B) Lambda (C) RDS

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

Answer: A

Explanation:

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

NEW QUESTION 75

If I want CloudFormation stack status updates to show up in a continuous delivery system in as close to real time as possible, how should I achieve this?

- A. Use a long-poll on the Resources object in your CloudFormation stack and display those state changes in the UI for the system.
- B. Use a long-poll on the `ListStacks` API call for your CloudFormation stack and display those state changes in the UI for the system.
- C. Subscribe your continuous delivery system to an SNS topic that you also tell your CloudFormation stack to publish events into.
- D. Subscribe your continuous delivery system to an SQS queue that you also tell your CloudFormation stack to publish events into.

Answer: C

Explanation:

Use NotificationARNs.member.N when making a CreateStack call to push stack events into SNS in nearly real-time.

Reference:

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

NEW QUESTION 77

What does it mean if you have zero IOPS and a non-empty I/O queue for all EBS volumes attached to a running EC2 instance?

- A. The I/O queue is buffer flushing.
- B. Your EBS disk head(s) is/are seeking magnetic stripes.
- C. The EBS volume is unavailable.
- D. You need to re-mount the EBS volume in the OS

Answer: C

Explanation:

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

"Unavailable" and "Unavailability" mean... For Amazon EBS, when all of your attached volumes perform zero read write IO, with pending IO in the queue.

Reference: <https://aws.amazon.com/ec2/sla/>

NEW QUESTION 81

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

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

Answer: D

Explanation:

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

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

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

NEW QUESTION 82

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 85

There is a very serious outage at AWS. EC2 is not affected, but your EC2 instance deployment scripts stopped working in the region with the outage. What might be the issue?

- A. The AWS Console is down, so your CLI commands do not work.
- B. S3 is unavailable, so you can't create EBS volumes from a snapshot you use to deploy new volumes.
- C. AWS turns off the `DeployCode` API call when there are major outages, to protect from system floods.
- D. None of the other answers make sense
- E. If EC2 is not affected, it must be some other issue

Answer: B

Explanation:

S3 stores all snapshots. If S3 is unavailable, snapshots are unavailable.

Amazon EC2 also uses Amazon S3 to store snapshots (backup copies) of the data volumes. You can use snapshots for recovering data quickly and reliably in case of application or system failures. You can also use snapshots as a baseline to create multiple new data volumes, expand the size of an existing data volume, or move data volumes across multiple Availability Zones, thereby making your data usage highly scalable. For more information about using data volumes and snapshots, see Amazon Elastic Block Store.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonS3.html>

NEW QUESTION 89

What is a circular dependency in AWS CloudFormation?

- A. When a Template references an earlier version of itself.
- B. When Nested Stacks depend on each other.
- C. When Resources form a DependOn loop.
- D. When a Template references a region, which references the original Template

Answer: C

Explanation:

To resolve a dependency error, add a DependsOn attribute to resources that depend on other resources in your template. In some cases, you must explicitly declare dependencies so that AWS CloudFormation can create or delete resources in the correct order. For example, if you create an Elastic IP and a VPC with an Internet gateway in the same stack, the Elastic IP must depend on the Internet gateway attachment. For additional information, see DependsOn Attribute.

Reference: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/troubleshooting.html#troubleshooting-g-errors-dependence-error>

NEW QUESTION 90

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 95

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 99

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 102

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

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

Answer: D

Explanation:

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

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

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

NEW QUESTION 103

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 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 106

When thinking of AWS Elastic Beanstalk, which statement is true?

- A. Worker tiers pull jobs from SNS.
- B. Worker tiers pull jobs from HTTP.
- C. Worker tiers pull jobs from JSON.
- D. Worker tiers pull jobs from SQ

Answer: D

Explanation:

Elastic Beanstalk installs a daemon on each Amazon EC2 instance in the Auto Scaling group to process Amazon SQS messages in the worker environment. The daemon pulls data off the Amazon SQS queue, inserts it into the message body of an HTTP POST request, and sends it to a user-configurable URL path on the local host. The content type for the message body within an HTTP POST request is application/json by default.

Reference:

<http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-tiers.html>

NEW QUESTION 109

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 114

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

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

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

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