



Amazon

Exam Questions AWS-Certified-DevOps-Engineer-Professional

Amazon AWS Certified DevOps Engineer Professional

NEW QUESTION 1

What method should I use to author automation if I want to wait for a CloudFormation stack to finish completing in a script?

- A. Event subscription using SQS.
- B. Event subscription using SNS.
- C. Poll using `ListStacks` / `list-stacks`.
- D. Poll using `GetStackStatus` / `get-stack-status`.

Answer: C

Explanation:

Event driven systems are good for IFTTT logic, but only polling will make a script wait to complete. ListStacks / list-stacks is a real method, GetStackStatus / get-stack-status is not.

Reference: <http://docs.aws.amazon.com/cli/latest/reference/cloudformation/list-stacks.html>

NEW QUESTION 2

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 3

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

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

Answer: B

Explanation:

Since access and storage for games is uniform, and you need to have ordering within each game for the scores (to access the highest value), your hash (partition) key should be the GameID, and there should be a range key for HighestScore.

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

NEW QUESTION 4

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 5

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 6

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 7

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

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

Answer: A

Explanation:

You can create a health check that will run into perpetuity using Route53, in one API call, which will ping your service via HTTP every 10 or 30 seconds.

Amazon Route 53 must be able to establish a TCP connection with the endpoint within four seconds. In addition, the endpoint must respond with an HTTP status code of 200 or greater and less than 400 within two seconds after connecting.

Reference:

<http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover-determining-health-of-endpoints.html>

NEW QUESTION 8

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 9

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 10

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 10

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

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

Answer: C

Explanation:

Even if you're careful to stay within the free tier, it's a good idea to create a billing alarm to notify you if you exceed the limits of the free tier. Billing alarms can help to protect you against unknowingly accruing charges if you inadvertently use a service outside of the free tier or if traffic exceeds your expectations. Reference:

<http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/free-tier-alarms.html>

NEW QUESTION 13

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-update-rollback.html>

NEW QUESTION 17

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 19

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/awsccloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html>

NEW QUESTION 23

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 24

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 28

You are experiencing performance issues writing to a DynamoDB table. Your system tracks high scores for video games on a marketplace. Your most popular game experiences all of the performance issues. What is the most likely problem?

- A. DynamoDB's vector clock is out of sync, because of the rapid growth in request for the most popular game.
- B. You selected the Game ID or equivalent identifier as the primary partition key for the table.
- C. Users of the most popular video game each perform more read and write requests than average.
- D. You did not provision enough read or write throughput to the table.

Answer: B

Explanation:

The primary key selection dramatically affects performance consistency when reading or writing to DynamoDB. By selecting a key that is tied to the identity of the game, you forced DynamoDB to create a hotspot in the table partitions, and over-request against the primary key partition for the popular game. When it stores data, DynamoDB divides a table's items into multiple partitions, and distributes the data primarily based upon the partition key value. The provisioned throughput associated with a table is also divided evenly among the partitions, with no sharing of provisioned throughput across partitions. Reference:

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

NEW QUESTION 30

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

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

Answer: D

Explanation:

standard volumes, or Magnetic volumes, are best for: Cold workloads where data is infrequently accessed, or scenarios where the lowest storage cost is important.

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

NEW QUESTION 32

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 33

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 35

Which of these is not a CloudFormation Helper Script?

- A. cfn-signal
- B. cfn-hup
- C. cfn-request
- D. cfn-get-metadata

Answer: C

Explanation:

This is the complete list of CloudFormation Helper Scripts: cfn-init, cfn-signal, cfn-get-metadata, cfn-hup Reference: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-helper-scripts-reference.html>

NEW QUESTION 39

You need to replicate API calls across two systems in real time. What tool should you use as a buffer and transport mechanism for API call events?

- A. AWS SQS
- B. AWS Lambda
- C. AWS Kinesis
- D. AWS SNS

Answer: C

Explanation:

AWS Kinesis is an event stream service. Streams can act as buffers and transport across systems for in-order programmatic events, making it ideal for replicating API calls across systems.

A typical Amazon Kinesis Streams application reads data from an Amazon Kinesis stream as data records. These applications can use the Amazon Kinesis Client Library, and they can run on Amazon EC2 instances. The processed records can be sent to dashboards, used to generate alerts, dynamically change pricing and advertising strategies, or send data to a variety of other AWS services. For information about Streams features and pricing, see Amazon Kinesis Streams.

Reference: <http://docs.aws.amazon.com/kinesis/latest/dev/introduction.html>

NEW QUESTION 43

You are building a Ruby on Rails application for internal, non-production use which uses IV|ySQL as a database. You want developers without very much AWS experience to be able to deploy new code with a single command line push. You also want to set this up as simply as possible. Which tool is ideal for this setup?

- A. AWS CloudFormation
- B. AWS OpsWorks

- C. AWS ELB + EC2 with CLI Push
- D. AWS Elastic Beanstalk

Answer: D

Explanation:

Elastic Beanstalk's primary mode of operation exactly supports this use case out of the box. It is simpler than all the other options for this question. With Elastic Beanstalk, you can quickly deploy and manage applications in the AWS cloud without worrying about the infrastructure that runs those applications. AWS Elastic Beanstalk reduces management complexity without restricting choice or control. You simply upload your application, and Elastic Beanstalk automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring.
Reference: http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_Ruby_rails.html

NEW QUESTION 45

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 47

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

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

Answer: C

Explanation:

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

Reference:

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

NEW QUESTION 51

From a compliance and security perspective, which of these statements is true?

- A. You do not ever need to rotate access keys for AWS IAM Users.
- B. You do not ever need to rotate access keys for AWS IAM Roles, nor AWS IAM Users.
- C. None of the other statements are true.
- D. You do not ever need to rotate access keys for AWS IAM Role

Answer: D

Explanation:

IAM Role Access Keys are auto-rotated by AWS on your behalf; you do not need to rotate them.

The application is granted the permissions for the actions and resources that you've defined for the role through the security credentials associated with the role.

These security credentials are temporary and we

rotate them automatically. We make new credentials available at least five minutes prior to the expiration of the old credentials.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

NEW QUESTION 56

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

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

Answer: D

Explanation:

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

NEW QUESTION 58

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 63

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 64

When thinking of AWS OpsWorks, which of the following is true?

- A. Stacks have many layers, layers have many instances.
- B. Instances have many stacks, stacks have many layers.
- C. Layers have many stacks, stacks have many instances.
- D. Layers have many instances, instances have many stack

Answer: A

Explanation:

The stack is the core AWS OpsWorks component. It is basically a container for AWS resources—Amazon EC2 instances, Amazon RDS database instances, and so on—that have a common purpose and should be logically managed together. You define the stack's constituents by adding one or more layers. A layer represents a set of Amazon EC2 instances that serve a particular purpose, such as serving applications or hosting a database server. An instance represents a single computing resource, such as an Amazon EC2 instance.

Reference: <http://docs.aws.amazon.com/opsworks/latest/userguide/welcome.html>

NEW QUESTION 67

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 68

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 72

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 73

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 75

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 79

If you're trying to configure an AWS Elastic Beanstalk worker tier for easy debugging if there are problems finishing queue jobs, what should you configure?

- A. Configure Rolling Deployments.
- B. Configure Enhanced Health Reporting
- C. Configure Blue-Green Deployments.
- D. Configure a Dead Letter Queue

Answer: D

Explanation:

Elastic Beanstalk worker environments support Amazon Simple Queue Service (SQS) dead letter queues. A dead letter queue is a queue where other (source) queues can send messages that for some reason could not be successfully processed. A primary benefit of using a dead letter queue is the ability to sideline and isolate the unsuccessfully processed messages. You can then analyze any messages sent to the dead letter queue to try to determine why they were not successfully processed. Reference: <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-tiers.html#worker-deadletter>

NEW QUESTION 83

You have a high security requirement for your AWS accounts. What is the most rapid and sophisticated setup you can use to react to AWS API calls to your account?

- A. Subscription to AWS Config via an SNS Topic
- B. Use a Lambda Function to perform in-flight analysis and react to changes as they occur.
- C. Global AWS CloudTrail setup delivering to S3 with an SNS subscription to the deliver notifications, pushing into a Lambda, which inserts records into an ELK stack for analysis.
- D. Use a CloudWatch Rule ScheduleExpression to periodically analyze IAM credential log
- E. Push the deltas for events into an ELK stack and perform ad-hoc analysis there.
- F. CloudWatch Events Rules which trigger based on all AWS API calls, submitting all events to an AWS Kinesis Stream for arbitrary downstream analysis.

Answer: D

Explanation:

CloudWatch Events allow subscription to AWS API calls, and direction of these events into Kinesis Streams. This allows a unified, near real-time stream for all API calls, which can be analyzed with any tool(s) of your choosing downstream.

Reference: http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/EventTypes.html#api_event_type

NEW QUESTION 88

.....

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

What method should I use to author automation if I want to wait for a CloudFormation stack to finish completing in a script?

- A. Event subscription using SQS.
- B. Event subscription using SNS.
- C. Poll using `ListStacks` / `list-stacks`.
- D. Poll using `GetStackStatus` / `get-stack-status`.

Answer: C

Explanation:

Event driven systems are good for IFTTT logic, but only polling will make a script wait to complete. ListStacks / list-stacks is a real method, GetStackStatus / get-stack-status is not.

Reference: <http://docs.aws.amazon.com/cli/latest/reference/cloudformation/list-stacks.html>

NEW QUESTION 2

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 3

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

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

Answer: B

Explanation:

Since access and storage for games is uniform, and you need to have ordering within each game for the scores (to access the highest value), your hash (partition) key should be the GameID, and there should be a range key for HighestScore.

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

NEW QUESTION 4

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 5

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 6

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 7

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

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

Answer: A

Explanation:

You can create a health check that will run into perpetuity using Route53, in one API call, which will ping your service via HTTP every 10 or 30 seconds.

Amazon Route 53 must be able to establish a TCP connection with the endpoint within four seconds. In addition, the endpoint must respond with an HTTP status code of 200 or greater and less than 400 within two seconds after connecting.

Reference:

<http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover-determining-health-of-endpoints.html>

NEW QUESTION 8

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 9

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 10

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 10

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

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

Answer: C

Explanation:

Even if you're careful to stay within the free tier, it's a good idea to create a billing alarm to notify you if you exceed the limits of the free tier. Billing alarms can help to protect you against unknowingly accruing charges if you inadvertently use a service outside of the free tier or if traffic exceeds your expectations. Reference:

<http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/free-tier-alarms.html>

NEW QUESTION 13

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-continueupdaterollback.html>

NEW QUESTION 17

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 19

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/awsccloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html>

NEW QUESTION 23

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 24

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 28

You are experiencing performance issues writing to a DynamoDB table. Your system tracks high scores for video games on a marketplace. Your most popular game experiences all of the performance issues. What is the most likely problem?

- A. DynamoDB's vector clock is out of sync, because of the rapid growth in request for the most popular game.
- B. You selected the Game ID or equivalent identifier as the primary partition key for the table.
- C. Users of the most popular video game each perform more read and write requests than average.
- D. You did not provision enough read or write throughput to the table.

Answer: B

Explanation:

The primary key selection dramatically affects performance consistency when reading or writing to DynamoDB. By selecting a key that is tied to the identity of the game, you forced DynamoDB to create a hotspot in the table partitions, and over-request against the primary key partition for the popular game. When it stores data, DynamoDB divides a table's items into multiple partitions, and distributes the data primarily based upon the partition key value. The provisioned throughput associated with a table is also divided evenly among the partitions, with no sharing of provisioned throughput across partitions. Reference:

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

NEW QUESTION 30

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

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

Answer: D

Explanation:

standard volumes, or Magnetic volumes, are best for: Cold workloads where data is infrequently accessed, or scenarios where the lowest storage cost is important.

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

NEW QUESTION 32

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 33

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 35

Which of these is not a CloudFormation Helper Script?

- A. cfn-signal
- B. cfn-hup
- C. cfn-request
- D. cfn-get-metadata

Answer: C

Explanation:

This is the complete list of CloudFormation Helper Scripts: cfn-init, cfn-signal, cfn-get-metadata, cfn-hup Reference: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-helper-scripts-reference.html>

NEW QUESTION 39

You need to replicate API calls across two systems in real time. What tool should you use as a buffer and transport mechanism for API call events?

- A. AWS SQS
- B. AWS Lambda
- C. AWS Kinesis
- D. AWS SNS

Answer: C

Explanation:

AWS Kinesis is an event stream service. Streams can act as buffers and transport across systems for in-order programmatic events, making it ideal for replicating API calls across systems.

A typical Amazon Kinesis Streams application reads data from an Amazon Kinesis stream as data records. These applications can use the Amazon Kinesis Client Library, and they can run on Amazon EC2 instances. The processed records can be sent to dashboards, used to generate alerts, dynamically change pricing and advertising strategies, or send data to a variety of other AWS services. For information about Streams features and pricing, see Amazon Kinesis Streams.

Reference: <http://docs.aws.amazon.com/kinesis/latest/dev/introduction.html>

NEW QUESTION 43

You are building a Ruby on Rails application for internal, non-production use which uses IV|ySQL as a database. You want developers without very much AWS experience to be able to deploy new code with a single command line push. You also want to set this up as simply as possible. Which tool is ideal for this setup?

- A. AWS CloudFormation
- B. AWS OpsWorks

- C. AWS ELB + EC2 with CLI Push
- D. AWS Elastic Beanstalk

Answer: D

Explanation:

Elastic Beanstalk's primary mode of operation exactly supports this use case out of the box. It is simpler than all the other options for this question. With Elastic Beanstalk, you can quickly deploy and manage applications in the AWS cloud without worrying about the infrastructure that runs those applications. AWS Elastic Beanstalk reduces management complexity without restricting choice or control. You simply upload your application, and Elastic Beanstalk automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring.
Reference: http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_Ruby_rails.html

NEW QUESTION 45

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 47

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

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

Answer: C

Explanation:

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

Reference:

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

NEW QUESTION 51

From a compliance and security perspective, which of these statements is true?

- A. You do not ever need to rotate access keys for AWS IAM Users.
- B. You do not ever need to rotate access keys for AWS IAM Roles, nor AWS IAM Users.
- C. None of the other statements are true.
- D. You do not ever need to rotate access keys for AWS IAM Role

Answer: D

Explanation:

IAM Role Access Keys are auto-rotated by AWS on your behalf; you do not need to rotate them.

The application is granted the permissions for the actions and resources that you've defined for the role through the security credentials associated with the role.

These security credentials are temporary and we

rotate them automatically. We make new credentials available at least five minutes prior to the expiration of the old credentials.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

NEW QUESTION 56

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

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

Answer: D

Explanation:

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

NEW QUESTION 58

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 63

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 64

When thinking of AWS OpsWorks, which of the following is true?

- A. Stacks have many layers, layers have many instances.
- B. Instances have many stacks, stacks have many layers.
- C. Layers have many stacks, stacks have many instances.
- D. Layers have many instances, instances have many stack

Answer: A

Explanation:

The stack is the core AWS OpsWorks component. It is basically a container for AWS resources—Amazon EC2 instances, Amazon RDS database instances, and so on—that have a common purpose and should be logically managed together. You define the stack's constituents by adding one or more layers. A layer represents a set of Amazon EC2 instances that serve a particular purpose, such as serving applications or hosting a database server. An instance represents a single computing resource, such as an Amazon EC2 instance.

Reference: <http://docs.aws.amazon.com/opsworks/latest/userguide/welcome.html>

NEW QUESTION 67

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 68

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 72

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 73

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 75

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 don't 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 79

If you're trying to configure an AWS Elastic Beanstalk worker tier for easy debugging if there are problems finishing queue jobs, what should you configure?

- A. Configure Rolling Deployments.
- B. Configure Enhanced Health Reporting
- C. Configure Blue-Green Deployments.
- D. Configure a Dead Letter Queue

Answer: D

Explanation:

Elastic Beanstalk worker environments support Amazon Simple Queue Service (SQS) dead letter queues. A dead letter queue is a queue where other (source) queues can send messages that for some reason could not be successfully processed. A primary benefit of using a dead letter queue is the ability to sideline and isolate the unsuccessfully processed messages. You can then analyze any messages sent to the dead letter queue to try to determine why they were not successfully processed. Reference: <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-tiers.html#worker-deadletter>

NEW QUESTION 83

You have a high security requirement for your AWS accounts. What is the most rapid and sophisticated setup you can use to react to AWS API calls to your account?

- A. Subscription to AWS Config via an SNS Topic
- B. Use a Lambda Function to perform in-flight analysis and react to changes as they occur.
- C. Global AWS CloudTrail setup delivering to S3 with an SNS subscription to the deliver notifications, pushing into a Lambda, which inserts records into an ELK stack for analysis.
- D. Use a CloudWatch Rule ScheduleExpression to periodically analyze IAM credential log
- E. Push the deltas for events into an ELK stack and perform ad-hoc analysis there.
- F. CloudWatch Events Rules which trigger based on all AWS API calls, submitting all events to an AWS Kinesis Stream for arbitrary downstream analysis.

Answer: D

Explanation:

CloudWatch Events allow subscription to AWS API calls, and direction of these events into Kinesis Streams. This allows a unified, near real-time stream for all API calls, which can be analyzed with any tool(s) of your choosing downstream.

Reference: http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/EventTypes.html#api_event_type

NEW QUESTION 88

.....

Relate Links

100% Pass Your AWS-Certified-DevOps-Engineer-Professional Exam with Examible 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/>