

➤ **Vendor: Amazon**

➤ **Exam Code: AWS Certified DevOps Engineer - Professional**

➤ **Exam Name: Amazon AWS Certified DevOps Engineer - Professional**

➤ **Question 21 – Question 40**

[Visit PassLeader and Download Full Version AWS Certified DevOps Engineer - Professional Exam Dumps](#)

QUESTION 21

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.

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

QUESTION 22

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.

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

QUESTION 23

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
<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-helper-scripts-reference.html>

QUESTION 24

Your team wants to begin practicing continuous delivery using CloudFormation, to enable automated builds and deploys of whole, versioned stacks or stack layers. You have a 3-tier, mission-critical system. Which of the following is NOT a best practice for using CloudFormation in a continuous delivery environment?

- A. Use the AWS CloudFormation `ValidateTemplate` call before publishing changes to AWS.
- B. Model your stack in one template, so you can leverage CloudFormation's state management and dependency resolution to propagate all changes.
- C. Use CloudFormation to create brand new infrastructure for all stateless resources on each push, and run integration tests on that set of infrastructure.
- D. Parametrize the template and use `Mappings` to ensure your template works in multiple Regions.

Answer: B

Explanation:

Putting all resources in one stack is a bad idea, since different tiers have different life cycles and frequencies of change. For additional guidance about organizing your stacks, you can use two common frameworks: a multi-layered architecture and service-oriented architecture (SOA).

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html#organizingstacks>

QUESTION 25

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.

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

QUESTION 26

You are building a Ruby on Rails application for internal, non-production use which uses MySQL 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.

http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_Ruby_rails.html

QUESTION 27

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.

<https://aws.amazon.com/iam/faqs/>

QUESTION 28

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 Accounts. Once they are logged in, the secret token passed to that user is used to directly access resources on AWS, like AWS S3.
- B. Use JWT or SAML compliant systems to build authorization policies. Users log in with a username and password, and are given a token they can use indefinitely to make calls against the photo infrastructure.
- C. Use AWS API Gateway with a constantly rotating API Key to allow access from the client-side. Construct a custom build of the SDK and include S3 access in it.
- D. Create an AWS OAuth Service Domain and grant public signup and access to the domain. 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.

<https://blogs.aws.amazon.com/security/post/Tx3SYCORF5EKRC0/How-Does-Amazon-Cognito-Relate-to-Existing-Web-Identity-Federatio>

QUESTION 29

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

<https://aws.amazon.com/cloudtrail/faqs/>

QUESTION 30

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.

<https://aws.amazon.com/lambda/faqs/>

QUESTION 31

Which is not a restriction on AWS EBS Snapshots?

- A. Snapshots which are shared cannot be used as a basis for other snapshots.
- B. You cannot share a snapshot containing an AWS Access Key ID or AWS Secret Access Key.
- C. You cannot share unencrypted snapshots.
- D. Snapshot restorations are restricted to the region in which the snapshots are created.

Answer: A

Explanation:

Snapshots shared with other users are usable in full by the recipient, including but limited to the ability to base modified volumes and snapshots.

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

QUESTION 32

You need to deploy a new application version to production. Because the deployment is high-risk, you need to roll the new version out to users over a number of hours, to make sure everything is working correctly. You need to be able to control the proportion of users seeing the new version of the application down to the percentage point. You use ELB and EC2 with Auto Scaling Groups and custom AMIs with your code pre-installed assigned to Launch Configurations. There are no database-level changes during your deployment. You have been told you cannot spend too much money, so you must not increase the number of EC2 instances much at all during the deployment, but you also need to be able to switch back to the original version of code quickly if something goes wrong. What is the best way to meet these requirements?

- A. Create a second ELB, Auto Scaling Launch Configuration, and Auto Scaling Group using the Launch Configuration. Create AMIs with all code pre-installed. Assign the new AMI to the second Auto Scaling Launch Configuration. Use Route53 Weighted Round Robin Records to adjust the proportion of traffic hitting the two ELBs.
- B. Use the Blue-Green deployment method to enable the fastest possible rollback if needed. Create a full second stack of instances and cut the DNS over to the new stack of instances, and change the DNS back if a rollback is needed.
- C. Create AMIs with all code pre-installed. Assign the new AMI to the Auto Scaling Launch Configuration, to replace the old one. Gradually terminate instances running the old code (launched with the old Launch Configuration) and allow the new AMIs to boot to adjust the traffic balance to the new code. On rollback, reverse the process by doing the same thing, but changing the AMI on the Launch Config back to the original code.
- D. Migrate to use AWS Elastic Beanstalk. Use the established and well-tested Rolling Deployment setting AWS provides on the new Application Environment, publishing a zip bundle of the new code and adjusting the wait period to spread the deployment over time. Re-deploy the old code bundle to rollback if needed.

Answer: A

Explanation:

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

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

QUESTION 33

What is required to achieve gigabit network throughput on EC2? You already selected cluster-compute, 10GB instances with enhanced networking, and your workload is already network-bound, but you are not seeing 10 gigabit speeds.

- A. Enable bplex networking on your servers, so packets are non-blocking in both directions and there's no switching overhead.

- B. Ensure the instances are in different VPCs so you don't saturate the Internet Gateway on any one VPC.
- C. Select PIOPS for your drives and mount several, so you can provision sufficient disk throughput.
- D. Use a placement group for your instances so the instances are physically near each other in the same Availability Zone.

Answer: D

Explanation:

You are not guaranteed 10gigabit performance, except within a placement group. A placement group is a logical grouping of instances within a single Availability Zone. Using placement groups enables applications to participate in a low-latency, 10 Gbps network. Placement groups are recommended for applications that benefit from low network latency, high network throughput, or both.

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

QUESTION 34

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.

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

QUESTION 35

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.

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

QUESTION 36

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 Roles.

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.

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

QUESTION 37

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. <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.RDS.SecurityGroups.html>

QUESTION 38

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.

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.MultiAZ.html>

QUESTION 39

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

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

QUESTION 40

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

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

region ELB.

Answer: A

Explanation:

There is no such thing as a cross-region ELB, nor such thing as a cross-region Auto Scaling Group, nor such thing as a DynamoDB Multi-Region Table. The only option that makes sense is the cross-regional replication version with two ELBs and ASGs with Route53 Failover and Latency DNS.

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

[Visit PassLeader and Download Full Version AWS Certified DevOps Engineer - Professional Exam Dumps](#)