



Vendor: Amazon

Exam Code: AWS Certified Solutions Architect - Professional

Exam Name: AWS Certified Solutions Architect - Professional

Version: DEMO

QUESTION 1

You've been brought in as solutions architect to assist an enterprise customer with their migration of an e-commerce platform to Amazon Virtual Private Cloud (VPC).

The previous architect has already deployed a 3-tier VPC.

The configuration is as follows:

```
VPC: vpc-2f8bc447
IGW: igw-2d8bc445
NACL: ad-208bc448
Subnets and Route Tables:
Web servers: subnet-258bc44d
Application servers: subnet-248bc44c
Database servers: subnet-9189c6f9
Route Tables:
rtb-218bc449
rtb-238bc44b
Associations:
subnet-258bc44d : rtb-218bc449
subnet-248bc44c : rtb-238bc44b
subnet-9189c6f9 : rtb-238bc44b
```

You are now ready to begin deploying EC2 instances into the VPC. Web servers must have direct access to the Internet. Application and database servers cannot have direct access to the Internet. Which configuration below will allow you the ability to remotely administer your application and database servers, as well as allow these servers to retrieve updates from the Internet?

- A. Create a bastion and NAT instance in subnet-258bc44d, and add a route from rtb-238bc44b to the NAT instance.
- B. Add a route from rtb-238bc44b to igw-2d8bc445 and add a bastion and NAT instance within subnet-248bc44c.
- C. Create a bastion and NAT instance in subnet-248bc44c, and add a route from rtb-238bc44b to subnet-258bc44d.
- D. Create a bastion and NAT instance in subnet-258bc44d, add a route from rtb-238bc44b to igw-2d8bc445, and a new NACL that allows access between subnet-258bc44d and subnet-248bc44c.

Answer: A

Explanation:

Create NAT instance in public subnet which is web server subnet (subnet-258bc44d) and add route (rtb-238bc44b) from private subnet (database subnet-9189c6f9) to the public NAT one to retrieve the updates.

QUESTION 2

You are designing the network infrastructure for an application server in Amazon VPC. Users will access all the application instances from the Internet, as well as from an on-premises network.

The on-premises network is connected to your VPC over an AWS Direct Connect link.

How would you design routing to meet the above requirements?

- A. Configure a single routing table with a default route via the Internet gateway. Propagate a default route via BGP on the AWS Direct Connect customer router. Associate the routing table with all VPC subnets.
- B. Configure a single routing table with a default route via the Internet gateway. Propagate specific routes for the on-premises networks via BGP on the AWS Direct Connect

customer router.

Associate the routing table with all VPC subnets.

- C. Configure two routing tables: one that has a default route via the Internet gateway, and another that has a default route via the VPN gateway.
Associate both routing tables with each VPC subnet.
- D. Configure a single routing table with two default routes: one to the Internet via an Internet gateway, the other to the on-premises network via the VPN gateway.
Use this routing table across all subnets in your VPC.

Answer: B

QUESTION 3

Within the IAM service a GROUP is regarded as a:

- A. A collection of AWS accounts
- B. It's the group of EC2 machines that gain the permissions specified in the GROUP.
- C. There's no GROUP in IAM, but only USERS and RESOURCES.
- D. A collection of users.

Answer: D

Explanation:

Use groups to assign permissions to IAM users

Instead of defining permissions for individual IAM users, it's usually more convenient to create groups that relate to job functions (administrators, developers, accounting, etc.), define the relevant permissions for each group, and then assign IAM users to those groups. All the users in an IAM group inherit the permissions assigned to the group. That way, you can make changes for everyone in a group in just one place. As people move around in your company, you can simply change what IAM group their IAM user belongs to.

<http://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html#use-groups-for-permissions>

QUESTION 4

Amazon EC2 provides a repository of public data sets that can be seamlessly integrated into AWS cloud-based applications.

What is the monthly charge for using the public data sets?

- A. A 1 time charge of 10\$ for all the datasets.
- B. 1\$ per dataset per month
- C. 10\$ per month for all the datasets
- D. There is no charge for using the public data sets

Answer: D

QUESTION 5

In the Amazon RDS Oracle DB engine, the Database Diagnostic Pack and the Database Tuning Pack are only available with _____

- A. Oracle Standard Edition
- B. Oracle Express Edition
- C. Oracle Enterprise Edition

D. None of these

Answer: C

Explanation:

<https://www.pythian.com/blog/a-most-simple-cloud-is-amazon-rds-for-oracle-right-for-you/>

QUESTION 6

You have deployed a web application, targeting a global audience across multiple AWS Regions under the domain name example.com.

You decide to use Route53 Latency-Based Routing to serve web requests to users from the region closest to the user. To provide business continuity in the event of server downtime you configure weighted record sets associated with two web servers in separate Availability Zones per region. During a DR test you notice that when you disable all web servers in one of the regions Route53 does not automatically direct all users to the other region.

What could be happening? Choose 2 answers

- A. You did not set "Evaluate Target Health" to 'Yes' on the latency alias resource record set associated with example.com in the region where you disabled the servers
- B. The value of the weight associated with the latency alias resource record set in the region with the disabled servers is higher than the weight for the other region
- C. One of the two working web servers in the other region did not pass its HTTP health check
- D. Latency resource record sets cannot be used in combination with weighted resource record sets
- E. You did not setup an HTTP health check for one or more of the weighted resource record sets associated with the disabled web servers

Answer: AE

Explanation:

<http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover-complex-configs.html>

For both latency alias resource record sets, you set the value of "Evaluate Target Health" to Yes. You use the Evaluate Target Health setting for each latency alias resource record set to make Amazon Route 53 evaluate the health of the alias targets—the weighted resource record sets—and respond accordingly.

QUESTION 7

An international company has deployed a multi-tier web application that relies on DynamoDB in a single region. For regulatory reasons they need disaster recovery capability in a separate region with a Recovery Time Objective of 2 hours and a Recovery Point Objective of 24 hours. They should synchronize their data on a regular basis and be able to provision the web application rapidly using CloudFormation.

The objective is to minimize changes to the existing web application, control the throughput of DynamoDB used for the synchronization of data, and synchronize only the modified elements. Which design would you choose to meet these requirements?

- A. Use AWS Data Pipeline to schedule a DynamoDB cross region copy once a day, create a "LastUpdated" attribute in your DynamoDB table that would represent the timestamp of the last update and use it as a filter
- B. Use AWS Data Pipeline to schedule an export of the DynamoDB table to S3 in the current region once a day, then schedule another task Immediately after it that will import data from S3 to DynamoDB in the other region
- C. Use EMR and write a custom script to retrieve data from DynamoDB in the current region using a SCAN operation and push it to DynamoDB in the second region
- D. Send also each write into an SQS queue in the second region, use an auto-scaling group behind

the SQS queue to replay the write in the second region

Answer: B

Explanation:

Export of dynamo DB is incremental and it will amend the backup with latest changes.

QUESTION 8

Your company currently has a 2-tier web application running in an on-premises data center. You have experienced several infrastructure failures in the past few months resulting in significant financial losses. Your CIO is strongly considering moving the application to AWS. While working on achieving buy-in from the other company executives, he asks you to develop a disaster recovery plan to help improve business continuity in the short term. He specifies a target Recovery Time Objective (RTO) of 4 hours and a Recovery Point Objective (RPO) of 1 hour or less. He also asks you to implement the solution within 2 weeks. Your database is 200GB in size and you have a 20Mbps Internet connection. How would you do this while minimizing costs?

- A. Create an EBS backed private AMI which includes a fresh install of your application. Develop a CloudFormation template which includes your AMI and the required EC2, AutoScaling, and ELB resources to support deploying the application across Multiple-Availability-Zones. Asynchronously replicate transactions from your on-premises database to a database instance in AWS across a secure VPN connection.
- B. Deploy your application on EC2 instances within an Auto Scaling group across multiple availability zones. Asynchronously replicate transactions from your on-premises database to a database instance in AWS across a secure VPN connection.
- C. Create an EBS backed private AMI which includes a fresh install of your application. Setup a script in your data center to backup the local database every 1 hour and to encrypt and copy the resulting file to an S3 bucket using multi-part upload.
- D. Install your application on a compute-optimized EC2 instance capable of supporting the application's average load. Synchronously replicate transactions from your on-premises database to a database instance in AWS across a secure Direct Connect connection.

Answer: A

QUESTION 9

You would like to create a mirror image of your production environment in another region for disaster recovery purposes. Which of the following AWS resources do not need to be recreated in the second region? Choose 2 answers

- A. Route53 Record Sets
- B. Launch Configurations
- C. EC2 Key Pairs
- D. Security Groups
- E. IAM Roles
- F. Elastic IP Addresses (EIP)

Answer: AE

Explanation:

As per the document defined, new IPs should be reserved not the same ones.

Elastic IP Addresses are static IP addresses designed for dynamic cloud computing. Unlike traditional static IP addresses, however, Elastic IP addresses enable you to mask instance or Availability Zone failures by programmatically remapping your public IP addresses to instances in your account in a particular region. For DR, you can also pre-allocate some IP addresses for the most critical systems so that their IP addresses are already known before disaster strikes. This can simplify the execution of the DR plan.

QUESTION 10

Your startup wants to implement an order fulfillment process for selling a personalized gadget that needs an average of 3-4 days to produce with some orders taking up to 6 months. You expect 10 orders per day on your first day, 1000 orders per day after 6 months and 10,000 orders after 12 months. Orders coming in are checked for consistency, then dispatched to your manufacturing plant for production, quality control, packaging, shipment and payment processing. If the product does not meet the quality standards at any stage of the process, employees may force the process to repeat a step. Customers are notified via email about order status and any critical issues with their orders such as payment failure.

Your base architecture includes AWS Elastic Beanstalk for your website with an RDS MySQL instance for customer data and orders.

How can you implement the order fulfillment process while making sure that the emails are delivered reliably?

- A. Add a business process management application to your Elastic Beanstalk app servers and re-use the RDS database for tracking order status.
Use one of the Elastic Beanstalk instances to send emails to customers.
- B. Use SWF with an Auto Scaling group of activity workers and a decider instance in another Auto Scaling group with min/max=1.
Use SES to send emails to customers.
- C. Use an SQS queue to manage all process tasks. Use an Auto Scaling group of EC2 instances that poll the tasks and execute them.
Use SES to send emails to customers.
- D. Use SWF with an Auto Scaling group of activity workers and a decider instance in another Auto Scaling group with min/max=1.
Use the decider instance to send emails to customers.

Answer: B

Explanation:

http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_ecommerce_checkout_13.pdf

Thank You for Trying Our Product

Braindump2go Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.braindump2go.com/all-products.html>



10% Discount Coupon Code: BDN2014