

Exam Questions DVA-C01

AWS Certified Developer Associate Exam

<https://www.2passeasy.com/dumps/DVA-C01/>



NEW QUESTION 1

A user is running a MySQL RDS instance. The user will not use the DB for the next 3 months. How can the user save costs?

- A. Pause the RDS instance from CLI until it is required in the future
- B. Stop the RDS instance
- C. Create a snapshot of RDS to launch in the future and terminate the instance now
- D. Change the instance size to micro

Answer: C

Explanation:

The RDS instances unlike the AWS EBS backed instances cannot be stopped or paused. The user needs to take the final snapshot, terminate the instance and launch a new instance in the future from that snapshot

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

NEW QUESTION 2

You have been doing a lot of testing of your VPC Network by deliberately failing EC2 instances to test whether instances are failing over properly. Your customer who will be paying the AWS bill for all this asks you if he being charged for all these instances. You try to explain to him how the billing works on EC2 instances to the best of your knowledge. What would be an appropriate response to give to the customer in regards to this?

- A. Billing commences when Amazon EC2 AMI instance is completely up and billing ends as soon as the instance starts to shutdown.
- B. Billing commences when Amazon EC2 initiates the boot sequence of an AMI instance and billing ends when the instance shuts down.
- C. Billing only commences only after 1 hour of uptime and billing ends when the instance terminates.
- D. Billing commences when Amazon EC2 initiates the boot sequence of an AMI instance and billing ends as soon as the instance starts to shutdown.

Answer: B

Explanation:

Billing commences when Amazon EC2 initiates the boot sequence of an AMI instance. Billing ends when the instance shuts down, which could occur through a web services command, by running "shutdown -h", or through instance failure.

Reference: <http://aws.amazon.com/ec2/faqs/#Billing>

NEW QUESTION 3

A user has launched five instances with ELB. How can the user add the sixth EC2 instance to ELB?

- A. The user can add the sixth instance on the fly.
- B. The user must stop the ELB and add the sixth instance.
- C. The user can add the instance and change the ELB config file.
- D. The ELB can only have a maximum of five instance

Answer: A

Explanation:

Elastic Load Balancing automatically distributes incoming traffic across multiple EC2 instances. You create a load balancer and register instances with the load balancer in one or more Availability Zones. The load balancer serves as a single point of contact for clients. This enables you to increase the availability of your application. You can add and remove EC2 instances from your load balancer as your needs change, without disrupting the overall flow of information.

Reference: <http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/SvcIntro.html>

NEW QUESTION 4

An organization has 500 employees. The organization wants to set up AWS access for each department. Which of the below mentioned options is a possible solution?

- A. Create IAM roles based on the permission and assign users to each role
- B. Create IAM users and provide individual permission to each
- C. Create IAM groups based on the permission and assign IAM users to the groups
- D. It is not possible to manage more than 100 IAM users with AWS

Answer: C

Explanation:

An IAM group is a collection of IAM users. Groups let the user specify permissions for a collection of users, which can make it easier to manage the permissions for those users.

Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_WorkingWithGroupsAndUsers.html

NEW QUESTION 5

In regard to DynamoDB, which of the following statements is correct?

- A. An Item should have at least two value sets, a primary key and another attribute.
- B. An Item can have more than one attributes.
- C. A primary key should be single-valued.
- D. An attribute can have one or several other attribute

Answer: B

Explanation:

In Amazon DynamoDB, a database is a collection of tables. A table is a collection of items and each item is a collection of attributes.

Reference: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DataModel.html>

NEW QUESTION 6

Which one of the following statements is NOT an advantage of DyanamoDB being built on Solid State Drives:

- A. serve high-scale request workloads
- B. low request pricing
- C. high I/O performance of WebApp on EC2 instance
- D. low-latency response times

Answer: C

Explanation:

In DynamoDB, SSDs help achieve design goals of predictable low-latency response times for storing and accessing data at any scale. The high I/O performance of SSDs also enables to serve high-scale request workloads cost efficiently, and to pass this efficiency along in low request pricing.

Reference: <http://aws.amazon.com/dynamodb/faqs/>

NEW QUESTION 7

A user is planning to make a mobile game which can be played online or offline and will be hosted on EC2.

The user wants to ensure that if someone breaks the highest score or they achieve some milestone they can inform all their colleagues through email. Which of the below mentioned AWS services helps achieve this goal?

- A. AWS Simple Workflow Service.
- B. AWS Simple Queue Service.
- C. Amazon Cognito
- D. AWS Simple Email Servic

Answer: D

Explanation:

Amazon Simple Email Service (Amazon SES) is a highly scalable and cost-effective email-sending service for businesses and developers. It integrates with other AWS services, making it easy to send emails from applications that are hosted on AWS.

Reference: <http://aws.amazon.com/ses/faqs/>

NEW QUESTION 8

Which one of the following operations is NOT a DynamoDB operation?

- A. BatchWriteItem
- B. DescribeTable
- C. BatchGetItem
- D. BatchDeleteItem

Answer: D

Explanation:

In DynamoDB, DeleteItem deletes a single item in a table by primary key, but BatchDeleteItem doesn't exist.

Reference: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/operationlist.html>

NEW QUESTION 9

Regarding Amazon SNS, when you want to subscribe to a topic and receive notifications to your email, in the Protocol drop-down box, you should select .

- A. Email
- B. Message
- C. SMTP
- D. IMAP

Answer: A

Explanation:

In Amazon SNS, when you want to subscribe to a topic and receive notifications to your email, select Email in the Protocol drop-down box. Enter an email address you can use to receive the notification in the Endpoint field.

Reference: <http://docs.aws.amazon.com/sns/latest/dg/SubscribeTopic.html>

NEW QUESTION 10

In DynamoDB, to get a detailed listing of secondary indexes on a table, you can use the action.

- A. DescribeTable
- B. BatchGetItem
- C. GetItem
- D. TableName

Answer: A

Explanation:

In DynamoDB, DescribeTable returns information about the table, including the current status of the table, when it was created, the primary key schema, and any

indexes on the table.

Reference: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/SecondaryIndexes.html>

NEW QUESTION 10

When you create a table with a hash-and-range key, you must define one or more secondary indexes on that table.

- A. False, hash-range key is another name for secondary index
- B. False, it is optional
- C. True
- D. False, when you have Hash-Range key you cannot define Secondary index

Answer: B

Explanation:

When you create a table with a hash-and-range key in DynamoDB, you can also define one or more secondary indexes on that table.

Reference: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/LSI.html>

NEW QUESTION 15

A user is planning to host a scalable dynamic web application on AWS. Which of the services may not be required by the user to achieve automated scalability?

- A. CloudWatch
- B. S3
- C. AutoScaling
- D. AWS EC2 instances

Answer: B

Explanation:

The user can achieve automated scaling by launching different EC2 instances and making them a part of an ELB. Cloudwatch will be used to monitor the resources and based on the scaling need it will trigger policies. AutoScaling is then used to scale up or down the instances.

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

NEW QUESTION 20

A user plans to use RDS as a managed DB platform. Which of the below mentioned features is not supported by RDS?

- A. Automated backup
- B. Automated scaling to manage a higher load
- C. Automated failure detection and recovery
- D. Automated software patching

Answer: B

Explanation:

AWS RDS provides a managed DB platform, which offers features, such as automated backup, patch management, automated failure detection and recovery. The scaling is not automated and the user needs to plan it with a few clicks.

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

NEW QUESTION 25

What happens if your application performs more reads or writes than your provisioned capacity?

- A. Nothing
- B. requests above your provisioned capacity will be performed but you will receive 400 error codes.
- C. requests above your provisioned capacity will be performed but you will receive 200 error codes.
- D. requests above your provisioned capacity will be throttled and you will receive 400 error code

Answer: D

Explanation:

Speaking about DynamoDB, if your application performs more reads/second or writes/second than your table's provisioned throughput capacity allows, requests above your provisioned capacity will be throttled and you will receive 400 error codes.

Reference: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/ProvisionedThroughputIntro.html>

NEW QUESTION 26

In relation to Amazon SQS, how can you ensure that messages are delivered in order?

- A. Increase the size of your queue
- B. Send them with a timestamp
- C. Give each message a unique id.
- D. AWS cannot guarantee that you will receive messages in the exact order you sent them

Answer: D

Explanation:

Amazon SQS makes a best effort to preserve order in messages, but due to the distributed nature of the queue, AWS cannot guarantee that you will receive messages in the exact order you sent them. You typically place sequencing information or timestamps in your messages so that you can reorder them upon receipt.

Reference: <https://aws.amazon.com/items/1343?externalID=1343>

NEW QUESTION 27

An orgAMzation has launched two applications: one for blogging and one for ECM on the same AWS Linux EC2 instance running in the AWS VPC. The orgAMzation has attached two private IPs (primary and secondary) to the above mentioned instance. The orgAMzation wants the instance OS to recognize the secondary IP address. How can the orgAMzation configure this?

- A. Use the ec2-net-utility package which updates routing tables, uses DHCP to refresh the secondary IP and adds the network interface.
- B. Use the ec2-net-utils package which will configure an additional network interface and update the routing table
- C. Use the ec2-ip-update package which can configure the network interface as well as update the secondary IP with DHCP.
- D. Use the ec2-ip-utility package which can update the routing tables as well as refresh the secondary IP using DHCP.

Answer: B

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. With VPC the user can specify multiple private IP addresses for his instances. The number of network interfaces and private IP addresses that a user can specify for an instance depends on the instance type. This scenario helps when the user wants to host multiple websites on a single EC2 instance. After the user has assigned a secondary private IP address to his instance, he needs to configure the operating system on that instance to recognize the secondary private IP address. For AWS Linux, the ec2-net-utils package can take care of this step. It configures additional network interfaces that the user can attach while the instance is running, refreshes secondary IP addresses during DHCP lease renewal, and updates the related routing rules.

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

NEW QUESTION 28

What kind of service is provided by AWS DynamoDB?

- A. Relational Database
- B. NoSQL Database
- C. Dynamic Database
- D. Document Database

Answer: B

Explanation:

DynamoDB is a fast, fully managed NoSQL database service. Reference: <http://aws.amazon.com/dynamodb/>

NEW QUESTION 29

Doug has created a VPC with CIDR 10.201.0.0/16 in his AWS account. In this VPC he has created a public subnet with CIDR block 10.201.31.0/24. While launching a new EC2 from the console, he is not able to assign the private IP address 10.201.31.6 to this instance. Which is the most likely reason for this issue?

- A. Private IP address 10.201.31.6 is not part of the associated subnet's IP address range.
- B. Private IP address 10.201.31.6 is blocked via ACLs in Amazon infrastructure as a part of platform security.
- C. Private address IP 10.201.31.6 is currently assigned to another interface.
- D. Private IP address 10.201.31.6 is reserved by Amazon for IP networking purpose

Answer: C

Explanation:

In Amazon VPC, you can assign any Private IP address to your instance as long as it is: Part of the associated subnet's IP address range
Not reserved by Amazon for IP networking purposes
Not currently assigned to another interface
Reference: <http://aws.amazon.com/vpc/faqs/>

NEW QUESTION 30

Regarding Amazon SQS, are there restrictions on the names of Amazon SQS queues?

- A. No
- B. Yes
- C. Queue names must be unique within an AWS account and you cannot use hyphens (-) and underscores (_)
- D. Yes
- E. Queue names are limited to 80 characters and queue names must be unique within an AWS account
- F. Yes
- G. Queue names are limited to 80 characters but queue names do not need to be unique within an AWS account

Answer: C

Explanation:

Queue names are limited to 80 characters. Alphanumeric characters plus hyphens (-) and underscores (_) are allowed. Queue names must be unique within an AWS account. After you delete a queue, you can reuse the queue name.

Reference: <https://aws.amazon.com/sqs/faqs/>

NEW QUESTION 35

Regarding Amazon SNS, to begin using Amazon SNS mobile push notifications, you first need that uses one of the supported push notification services: APNS, GCM, or ADM.

- A. an access policy for the mobile endpoints
- B. to active push notification service of Amazon SNS
- C. to know the type of mobile device operating system
- D. an app for the mobile endpoints

Answer:

D

Explanation:

In Amazon SNS, to begin using Amazon SNS mobile push notifications, you first need an app for the mobile endpoints that uses one of the supported push notification services: APNS, GCM, or ADM. After you've registered and configured the app to use one of these services, you configure Amazon SNS to send push notifications to the mobile endpoints.

Reference: <http://docs.aws.amazon.com/sns/latest/dg/SNSMobilePush.html>

NEW QUESTION 36

ExamKiller (with AWS account ID H112223333) has created 50 IAM users for its organization's employees. ExamKiller wants to make the AWS console login URL for all IAM users as: <https://examkiller.signin.aws.amazon.com/console/>. How can this be configured?

- A. Create a bucket with the name ExamKiller and map it with the IAM alias
- B. It is not possible to have capital letters as a part of the alias name
- C. The user needs to use Route 53 to map the ExamKiller domain and IAM URL
- D. For the AWS account, create an alias ExamKiller for the IAM login

Answer: B

Explanation:

If a user wants the URL of the AWS IAM sign-in page to have the company name instead of the AWS account ID, he can create an alias for his AWS account ID. The alias must be unique across all Amazon WebServices products and contain only digits, lowercase letters, and hyphens.

Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/AccountAlias.html>

NEW QUESTION 40

Can a user get a notification of each instance start / terminate configured with Auto Scaling?

- A. Yes, always
- B. No
- C. Yes, if configured with the Auto Scaling group
- D. Yes, if configured with the Launch Config

Answer: C

Explanation:

The user can get notifications using SNS if he has configured the notifications while creating the Auto Scaling group.

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

NEW QUESTION 41

AutoScaling is configured with 3 AZs. Each zone has 5 instances running. If AutoScaling wants to terminate an instance based on the policy action, which instance will it terminate first?

- A. Terminate the first launched instance
- B. Randomly select the instance for termination
- C. Terminate the instance from the AZ which does not have a high AWS load
- D. Terminate the instance from the AZ which has instances running near to the billing hour

Answer: B

Explanation:

Before Auto Scaling selects an instance to terminate, it first identifies the Availability Zone that has more instances than the other Availability Zones used by the group. If all the Availability Zones have the same number of instances, it identifies a random Availability Zone.

Reference: <http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/us-termination-policy.html>

NEW QUESTION 45

In regard to DynamoDB, can I delete local secondary indexes?

- A. Yes, if it is a primary hash key index
- B. No
- C. Yes, if it is a local secondary indexes
- D. Yes, if it is a Global secondary indexes

Answer: B

Explanation:

In DynamoDB, an index cannot be modified once it is created. Reference: http://aws.amazon.com/dynamodb/faqs/#security_anchor

NEW QUESTION 50

You need to develop and run some new applications on AWS and you know that Elastic Beanstalk and CloudFormation can both help as a deployment mechanism for a broad range of AWS resources. Which of the following statements best describes the differences between Elastic Beanstalk and CloudFormation?

- A. Elastic Beanstalk uses Elastic load balancing and CloudFormation doesn't.
- B. CloudFormation is faster in deploying applications than Elastic Beanstalk.
- C. CloudFormation is much more powerful than Elastic Beanstalk, because you can actually design and script custom resources
- D. Elastic Beanstalk is faster in deploying applications than CloudFormation

Answer: C

Explanation:

These services are designed to complement each other. AWS Elastic Beanstalk provides an environment to easily develop and run applications in the cloud. It is integrated with developer tools and provides a one-stop experience for you to manage the lifecycle of your applications. AWS CloudFormation is a convenient deployment mechanism for a broad range of AWS resources. It supports the infrastructure needs of many different types of applications such as existing enterprise applications, legacy applications, applications built using a variety of AWS resources and container-based solutions (including those built using AWS Elastic Beanstalk).

AWS CloudFormation introduces two new concepts: The template, a JSON-format, text-based file that describes all the AWS resources you need to deploy to run your application and the stack, the set of AWS resources that are created and managed as a single unit when AWS CloudFormation instantiates a template.

Reference: <http://aws.amazon.com/cloudformation/faqs/>

NEW QUESTION 51

A user is planning to host a web server as well as an app server on a single EC2 instance which is a part of the public subnet of a VPC. How can the user setup to have two separate public IPs and separate security groups for both the application as well as the web server?

- A. Launch a VPC instance with two network interface
- B. Assign a separate security group to each and AWS will assign a separate public IP to them.
- C. Launch VPC with two separate subnets and make the instance a part of both the subnets.
- D. Launch a VPC instance with two network interface
- E. Assign a separate security group and elastic IP to them.
- F. Launch a VPC with ELB such that it redirects requests to separate VPC instances of the public subne

Answer: C

Explanation:

If you need to host multiple websites(with different IPs) on a single EC2 instance, the following is the suggested method from AWS.

Launch a VPC instance with two network interfaces

Assign elastic IPs from VPC EIP pool to those interfaces (Because, when the user has attached more than one network interface with an instance, AWS cannot assign public IPs to them.)

Assign separate Security Groups if separate Security Groups are needed

This scenario also helps for operating network appliances, such as firewalls or load balancers that have multiple private IP addresses for each network interface.

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

NEW QUESTION 53

How long are the messages kept on an SQS queue by default?

- A. If a message is not read, it is never deleted
- B. 2 weeks
- C. 1 day
- D. 4 days

Answer: D

Explanation:

The SQS message retention period is configurable and can be set anywhere from 1 minute to 2 weeks. The default is 4 days and once the message retention limit is reached your messages will be automatically deleted. The option for longer message retention provides greater flexibility to allow for longer intervals between message production and consumption.

Reference: <https://aws.amazon.com/sqs/faqs/>

NEW QUESTION 56

How can a user configure three termination policies for the AutoScaling group?

- A. Define multiple policies in random order
- B. Define multiple policies in the ordered list
- C. Keep updating the AutoScaling group with each policy
- D. The user cannot specify more than two policies for AutoScaling

Answer: B

Explanation:

To configure the Auto Scaling termination policy, the user can either specify any one of the policies as a standalone policy or list multiple policies in an ordered list. The policies are executed in the order that they are listed.

Reference: <http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/us-termination-policy.html>

NEW QUESTION 59

A user wants to configure AutoScaling which scales up when the CPU utilization is above 70% and scales down when the CPU utilization is below 30%. How can the user configure AutoScaling for the above mentioned condition?

- A. Use AutoScaling with a schedule
- B. Configure ELB to notify AutoScaling on load increase or decrease
- C. Use dynamic AutoScaling with a policy
- D. Use AutoScaling by manually modifying the desired capacity during a condition

Answer: C

Explanation:

The user can configure the AutoScaling group to automatically scale up and then scale down based on the specified conditions. To configure this, the user must

setup policies which will get triggered by the CloudWatch alarms.

Reference:

<http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/as-scale-based-on-demand.html>

NEW QUESTION 60

A user has created an application which sends data to a log file. The server hosting the log files can be unavailable due to any reason. The user wants to make it so that whenever the log server is up it should be receiving the messages. Which of the below mentioned AWS services helps achieve this functionality?

- A. AWS Simple Workflow
- B. AWS Simple Task Service
- C. AWS Simple Notification Service
- D. AWS Simple Queue Service

Answer: D

Explanation:

Amazon Simple Queue Service (SQS) is a fast, reliable, scalable, and fully managed message queuing service. SQS provides a simple and cost-effective way to decouple the components of an application. The user can use SQS to transmit any volume of data without losing messages or requiring other services to always be available. Using SQS, the application has to just send the data to SQS and SQS transmits it to the log file whenever it is available.

Reference: <http://aws.amazon.com/sqs/>

NEW QUESTION 62

In AWS Elastic Beanstalk, you can update your deployed application even while it is part of a running environment. For a Java application, you can also use to update your deployed application.

- A. the AWS Toolkit for Eclipse
- B. the AWS Toolkit for Visual Studio
- C. the AWS Toolkit for JVM
- D. the AWS Toolkit for Netbeans

Answer: A

Explanation:

In AWS Elastic Beanstalk, you can update your deployed application, even while it is part of a running environment. For a Java application, you can also use the AWS Toolkit for Eclipse to update your deployed application.

Reference: <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/GettingStarted.Walkthrough.html>

NEW QUESTION 67

What is the maximum time messages can be stored in SQS?

- A. 14 days
- B. one month
- C. 4 days
- D. 7 days

Answer: A

Explanation:

A message can be stored in the Simple Queue Service (SQS) from 1 minute up to a maximum of 14 days. Reference:

http://aws.amazon.com/sqs/faqs/#How_long_can_I_keep_my_messages_in_Amazon_SQS_queues

NEW QUESTION 72

A user is launching an AWS RDS instance with MySQL. The user wants to enable the Multi AZ feature. Which of the below mentioned parameters will not be allowed to configure by RDS?

- A. Availability Zone
- B. Region
- C. DB subnet group
- D. Database port

Answer: A

Explanation:

If the user is launching RDS with Multi AZ the user cannot provision the Availability Zone. RDS is launched automatically instead

Reference: <https://console.aws.amazon.com/rds/>

NEW QUESTION 75

You want to have multiple versions of your application running at the same time, with all versions launched via AWS Elastic Beanstalk. Is this possible?

- A. N
- B. However if you have 2 AWS accounts this can be done
- C. N
- D. AWS Elastic Beanstalk is not designed to support multiple running environments
- E. Ye
- F. AWS Elastic Beanstalk is designed to support a number of multiple running environments
- G. Ye
- H. However AWS Elastic Beanstalk is designed to support only 2 multiple running environments

Answer: C

Explanation:

AWS Elastic Beanstalk is designed to support multiple running environments. As an example you could have one for integration testing, one for pre-production, and one for production, with each environment independently configured and running on its own separate AWS resources.

Reference: <https://aws.amazon.com/elasticbeanstalk/faqs/>

NEW QUESTION 76

A user has launched one EC2 instance in the US West region. The user wants to access the RDS instance launched in the US East region from that EC2 instance. How can the user configure the access for that EC2 instance?

- A. It is not possible to access RDS of the US East region from the US West region
- B. Open the security group of the US West region in the RDS security group's ingress rule
- C. Configure the IP range of the US West region instance as the ingress security rule of RDS
- D. Create an IAM role which has access to RDS and launch an instance in the US West region with it

Answer: C

Explanation:

The user cannot authorize an Amazon EC2 security group if it is in a different AWS Region than the RDS DB instance. The user can authorize an IP range or specify an Amazon EC2 security group in the same region that refers to an IP address in another region.

Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithSecurityGroups.html

NEW QUESTION 78

In regard to DynamoDB, what is the Global secondary index?

- A. An index with a hash and range key that can be different from those on the table.
- B. An index that has the same range key as the table, but a different hash key
- C. An index that has the same hash key and range key as the table
- D. An index that has the same hash key as the table, but a different range key

Answer: A

Explanation:

Global secondary index - an index with a hash and range key that can be different from those on the table.

Reference: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DataModel.html>

NEW QUESTION 83

How does Amazon SQS allow multiple readers to access the same message queue without losing messages or processing them many times?

- A. By identifying a user by his unique id
- B. By using unique cryptography
- C. Amazon SQS queue has a configurable visibility timeout.
- D. Multiple readers can't access the same message queue

Answer: C

Explanation:

Every Amazon SQS queue has a configurable visibility timeout. For the designated amount of time after a message is read from a queue, it will not be visible to any other reader. As long as the amount of time that it takes to process the message is less than the visibility timeout, every message will be processed and deleted. In the event that the component processing the message fails or becomes unavailable, the message will again become visible to any component reading the queue once the visibility timeout ends. This allows you to have many components all reading messages from the same queue, with each working to process different messages.

Reference: <https://aws.amazon.com/sqs/faqs/>

NEW QUESTION 84

A user is trying to understand AWS SNS. To which of the below mentioned end points is SNS unable to send a notification?

- A. AWS SES
- B. Email JSON
- C. AWS SQS
- D. HTTP

Answer: A

Explanation:

Amazon Simple Notification Service (Amazon SNS) is a fast, filexible, and fully managed push messaging service. Amazon SNS can deliver notifications by SMS text message or email to the Amazon Simple Queue Service (SQS) queues or to any HTTP endpoint. The user can select one the following transports as part of the subscription requests: "HTTP", "HTTPS", "Email", "Email-JSON", "SQS", "and SMS".

Reference: <http://aws.amazon.com/sns/faqs/>

NEW QUESTION 87

Can one instance be registered with two ELBs in the same region?

- A. No
- B. Yes, provided both ELBs have the same health check configuration
- C. Yes, always

D. Yes, provided both ELBs are in the same AZ

Answer: C

Explanation:

Yes, it is possible to have one instance part of two separate ELBs, though both ELBs have different configurations. ELBs are never launched in specific zones.

Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/enable-disable-az.html>

NEW QUESTION 90

A user has created photo editing software and hosted it on EC2. The software accepts requests from the user about the photo format and resolution and sends a message to S3 to enhance the picture accordingly. Which of the below mentioned AWS services will help make a scalable software with the AWS infrastructure in this scenario?

- A. AWS Elastic Transcoder
- B. AWS Simple Notification Service
- C. AWS Simple Queue Service
- D. AWS Glacier

Answer: C

Explanation:

Amazon Simple Queue Service (SQS) is a fast, reliable, scalable, and fully managed message queuing service. SQS provides a simple and cost-effective way to decouple the components of an application. The user can configure SQS, which will decouple the call between the EC2 application and S3. Thus, the application does not keep waiting for S3 to provide the data.

Reference: <http://aws.amazon.com/sqs/faqs/>

NEW QUESTION 91

A user is planning to host a mobile game on EC2 which sends notifications to active users on either high score or the addition of new features. The user should get this notification when he is online on his mobile device. Which of the below mentioned AWS services can help achieve this functionality?

- A. AWS Simple Notification Service.
- B. AWS Simple Queue Service.
- C. AWS Mobile Communication Service.
- D. AWS Simple Email Service

Answer: A

Explanation:

Amazon Simple Notification Service (Amazon SNS) is a fast, flexible, and fully managed push messaging service. Amazon SNS makes it simple and cost-effective to push to mobile devices, such as iPhone, iPad, Android, Kindle Fire, and internet connected smart devices, as well as pushing to other distributed services.

Reference: <http://aws.amazon.com/sns>

NEW QUESTION 93

An organization is setting up their website on AWS. The organization is working on various security measures to be performed on the AWS EC2 instances. Which of the below mentioned security mechanisms will not help the organization to avoid future data leaks and identify security weaknesses?

- A. Perform SQL injection for application testing.
- B. Run penetration testing on AWS with prior approval from Amazon.
- C. Perform a hardening test on the AWS instance.
- D. Perform a Code Check for any memory leak

Answer: D

Explanation:

AWS security follows the shared security model where the user is as much responsible as Amazon. Since Amazon is a public cloud it is bound to be targeted by hackers. If an organization is planning to host their application on AWS EC2, they should perform the below mentioned security checks as a measure to find any security weakness/data leaks:

Perform penetration testing as performed by attackers to find any vulnerability. The organization must take an approval from AWS before performing penetration testing

Perform hardening testing to find if there are any unnecessary ports open Perform SQL injection to find any DB security issues

The code memory checks are generally useful when the organization wants to improve the application performance.

Reference: <http://aws.amazon.com/security/penetration-testing/>

NEW QUESTION 94

A root account owner is trying to setup an additional level of security for all his IAM users. Which of the below mentioned options is a recommended solution for the account owner?

- A. Enable access key and secret access key for all the IAM users
- B. Enable MFA for all IAM users
- C. Enable the password for all the IAM users
- D. Enable MFA for the root account

Answer: B

Explanation:

Multi-Factor Authentication adds an extra level of security for all the users. The user can enable MFA for all IAM users which ensures that each user has to provide an extra six digit code for authentication. Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_ManagingMFA.html

NEW QUESTION 96

Regarding Amazon SQS, what happens if there is no actMty against a queue for more than 30 consecutive days?

- A. Your account will be suspended
- B. The queue may be deleted
- C. Nothing
- D. The queue will be deleted

Answer: B

Explanation:

AWS reserve the right to delete a queue if none of the following requests have been issued against the queue for more than 30 consecutive days:
SendMessage ReceiveMessage DeleteMessage GetQueueAttributes SetQueueAttributes
You should design your application with this in mind. Reference: <https://aws.amazon.com/sqs/faqs/>

NEW QUESTION 97

ExamKiller (with AWS account ID H1122223333) has created 50 IAM users for its orgAMzation's employees. ExamKil|er wants to make the AWS console login URL for all IAM users like: <https://examki|er.signin.aws.amazon.com/console/>. How can this be configured?

- A. The user needs to use Route 53 to map the examkiller domain and IAM URL
- B. Create an IAM AWS account alias with the name examkiller
- C. It is not possible to have a personalized IAM login URL
- D. Create an IAM hosted zone Identity for the domain examkiller

Answer: B

Explanation:

If a user wants the URL of the AWS IAM sign-in page to have a company name instead of the AWS account ID, he can create an alias for his AWS account ID.
Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/AccountAlias.html>

NEW QUESTION 102

A user has setup Multi AZ with the MS SQL RDS instance. Which of the below mentioned functionalities can be achieved by the user?

- A. High availability
- B. Scalability
- C. MS SQL does not support Multi AZ
- D. Disaster recovery

Answer: C

Explanation:

The Multi AZ feature allows the user to achieve High Availability. MS SQL does not support Multi AZ. Reference: <https://aws.amazon.com/rds/faqs/#36>

NEW QUESTION 107

A user has created an EBS instance in the US-East-1a AZ. The user has a volume of 30 GB in the US-East-1 b zone. How can the user attach the volume to an instance?

- A. Since both the volume and the instance are in the same region, the user can attach the volume
- B. Use the volume migrate function to move the volume from one AZ to another and attach to the instance
- C. Take a snapshot of the volum
- D. Create a new volume in the USEast-1a and attach that to the instance
- E. Use the volume replicate function to create a new volume in the US-East-1a and attach that to the volume

Answer: C

Explanation:

If an EBS volume is not in the same AZ of an EC2 instance, it cannot be attached to the instance. The only option is to take a snapshot of the volume and create a new volume in the instance's AZ. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSSnapshots.html>

NEW QUESTION 108

A user is part of a group which has a policy allowing him just read only access to EC2. The user is part of another group which has full access to EC2. What happens when the user tries to launch an instance?

- A. It will allow the user to launch the instance
- B. It will fail since the user has just read only access
- C. It will allow or deny based on the group under which the user has logged into EC2
- D. It will not allow the user to add to the conflicting groups

Answer: A

Explanation:

The IAM group policy is always aggregated. In this case, if the user does not have permission for one group, but has permission for another group, he will have full access to EC2. Unless there is specific deny policy, the user will be able to access EC2.
Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/PoliciesOverview.html>

NEW QUESTION 110

A user has launched an RDS instance. The user has created 3 databases on the same server. What can the maximum size be for each database?

- A. The size of each DB cannot be more than 3 TB
- B. It is not possible to have more than one DB on a single instance
- C. The total instance storage size cannot be more than 3 TB
- D. The size of each DB cannot be more than 1 TB

Answer: C

Explanation:

The AWS RDS DB instance is an isolated DB environment provided by AWS in which the user can create more than 1 database. The maximum size of the instance should be between 5 GB and 3 TB. The size of each DB can be anything in this range.

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

NEW QUESTION 114

A user is creating a new EBS volume from an existing snapshot. The snapshot size shows 10 GB. Can the user create a volume of 30 GB from that snapshot?

- A. Provided the original volume has set the change size attribute to true
- B. Yes
- C. Provided the snapshot has the modify size attribute set as true
- D. No

Answer: B

Explanation:

A user can always create a new EBS volume of a higher size than the original snapshot size. The user cannot create a volume of a lower size. When the new volume is created the size in the instance will be shown as the original size. The user needs to change the size of the device with `resize2fs` or other OS specific commands.

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

NEW QUESTION 118

can be used to bootstrap both the Chef Server and Chef Client software on your EC2 instances.

- A. AWS CloudFormation
- B. AWS Elastic Beanstalk
- C. AWS OpsWorks
- D. Amazon Glacier

Answer: A

Explanation:

AWS CloudFormation can be used to bootstrap both the Chef Server and Chef Client software on your EC2 instances.

Reference: <http://aws.amazon.com/cloudformation/faqs/>

NEW QUESTION 121

An organization has 20 employees. The organization wants to give all the users access to the organization AWS account. Which of the below mentioned options is the right solution?

- A. Share the root credentials with all the users
- B. Create an IAM user for each employee and provide access to them
- C. It is not advisable to give AWS access to so many users
- D. Use the IAM role to allow access based on STS

Answer: B

Explanation:

AWS Identity and Access Management is a web service that enables the AWS customers to manage users and user permissions in AWS. The IAM is targeted at organizations with multiple users or systems that use AWS products such as Amazon EC2, Amazon RDS, and the AWS Management Console. With IAM, the organization can centrally manage users, security credentials such as access keys, and permissions that control which AWS resources users can access.

Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/IAM_Introduction.htm

NEW QUESTION 125

When AutoScaling is launching a new instance based on condition, which of the below mentioned policies will it follow?

- A. Based on the criteria defined with cross zone Load balancing
- B. Launch an instance which has the highest load distribution
- C. Launch an instance in the AZ with the fewest instances
- D. Launch an instance in the AZ which has the highest instances

Answer: C

Explanation:

AutoScaling attempts to distribute instances evenly between the Availability Zones that are enabled for the user's AutoScaling group. Auto Scaling does this by attempting to launch new instances in the Availability Zone with the fewest instances.

Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AS_Concepts.html

NEW QUESTION 126

In regards to Amazon SQS how can you secure the messages in your queues?

- A. You can't
- B. Amazon SQS uses either your Access Key ID or an X.509 certificate to authenticate your identity
- C. Through your IAM access keys
- D. Don't use root access

Answer: B

Explanation:

Authentication mechanisms are provided to ensure that messages stored in Amazon SQS queues are secured against unauthorized access. Only the AWS account owners can access the queues they create. Amazon SQS uses proven cryptographic methods to authenticate your identity, either through the use of your Access Key ID and request signature, or through the use of an X.509 certificate.

Reference: <https://aws.amazon.com/sqs/faqs/>

NEW QUESTION 131

Which Amazon service is not used by Elastic Beanstalk?

- A. Amazon S3
- B. Amazon ELB
- C. Auto scaling
- D. Amazon EMR

Answer: D

Explanation:

Elastic Beanstalk leverages AWS services such as Amazon Elastic Cloud Compute (Amazon EC2), Amazon Simple Storage Service (Amazon S3), Amazon Simple Notification Service (Amazon SNS), Elastic Load Balancing and Auto Scaling to deliver the same highly reliable, scalable, and cost-effective infrastructure that hundreds of thousands of businesses depend on today.

Reference: <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/welcome.html>

NEW QUESTION 133

Which of the below mentioned options can be a good use case for storing content in AWS RRS?

- A. Storing mission critical data Files
- B. Storing infrequently used log files
- C. Storing a video file which is not reproducible
- D. Storing image thumbnails

Answer: D

Explanation:

AWS RRS provides the same functionality as AWS S3, but at a cheaper rate. It is ideally suited for non-mission, critical applications, such as files which can be reproduced.

Reference: <http://docs.aws.amazon.com/AmazonS3/latest/dev/UsingRRS.html>

NEW QUESTION 138

Which header received at the EC2 instance identifies the port used by the client while requesting ELB?

- A. X-Forwarded-Proto
- B. X-Requested-Proto
- C. X-Forwarded-Port
- D. X-Requested-Port

Answer: C

Explanation:

The X-Forwarded-Port request header helps the user identify the port used by the client while sending a request to ELB.

Reference: <http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/TerminologyandKeyConcepts.html>

NEW QUESTION 143

A user is trying to create a policy for an IAM user from the AWS console. Which of the below mentioned options is not available to the user while configuring policy?

- A. Use policy generator to create policy
- B. Use custom policy to create policy
- C. Use policy simulator to create policy
- D. Assign No permission

Answer: C

Explanation:

When a user is trying to create a policy from the AWS console, it will have options such as create policy from templates or use a policy generator. The user can also define a custom policy or chose the option to have no permission. The policy simulator is not available in the console.

Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/IAMBestPractices.html>

NEW QUESTION 147

A user has an S3 object in the US Standard region with the content "color=red". The user updates the object with the content as "color=white". If the user tries to read the value 1 minute after it was uploaded, what will S3 return?

- A. It will return "color=white"
- B. It will return "color=red"
- C. It will return an error saying that the object was not found
- D. It may return either "color=red" or "color=white" i.
- E. any of the value

Answer: D

Explanation:

AWS S3 follows the eventual consistent model in the US Standard Region. Once the object is updated it may return the new value or the old value based on whether all the content is replicated across multiple servers until it becomes consistent (eventual).
Reference: <http://docs.aws.amazon.com/AmazonS3/latest/dev/Introduction.html>

NEW QUESTION 151

A user is creating a snapshot of an EBS volume. Which of the below statements is incorrect in relation to the creation of an EBS snapshot?

- A. Its incremental
- B. It can be used to launch a new instance
- C. It is stored in the same AZ as the volume
- D. It is a point in time backup of the EBS volume

Answer: C

Explanation:

The EBS snapshots are a point in time backup of the EBS volume. It is an incremental snapshot, but is always specific to the region and never specific to a single AZ.

Hence the statement "It is stored in the same AZ as the volume" is incorrect.

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

NEW QUESTION 152

You are using Amazon SQS and are getting a "Queue Deleted Recently" error. What is wrong?

- A. The message is too big
- B. You have incorrect permissions
- C. Another user has deleted the queue
- D. If you delete a queue, you need to wait for at least 60 seconds before creating a queue with the same name

Answer: D

Explanation:

If you delete a queue, you need to wait for at least 60 seconds before creating a queue with the same name. Please note that when you delete a queue, the deletion process takes up to 60 seconds. Requests you send to a recently deleted queue might succeed during the 60-second period. For example, a `SendMessage` request might succeed, but after 60 seconds the queue and that message you sent no longer exists.

Reference: <https://aws.amazon.com/items/1343?externalID=1343>

NEW QUESTION 156

What is the maximum number of tags that a user can assign to an EC2 instance?

- A. 50
- B. 10
- C. 5
- D. 25

Answer: B

Explanation:

To help manage EC2 instances as well as their usage in a better way, the user can tag the instances. The tags are metadata assigned by the user which consists of a key and a value. One resource can have a maximum of 10 tags.

Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html

NEW QUESTION 157

An organization is hosting an application as part of the free usage tier. The organization wants to create IAM users for each of its 150 employees and they may access AWS as part of free usage tier. What will you advise the organization?

- A. The IAM is not available as a part of the free usage tier
- B. Create IAM roles and give access based on role since it will not cost the user
- C. Do not create more than 100 users as it will cost the organization.
- D. Create IAM users for each employee as it does not cost

Answer: D

Explanation:

IAM is a free service. You can create as many IAM users or groups as desired free of cost. Reference:

http://docs.aws.amazon.com/IAM/latest/UserGuide/IAM_Introduction.html

NEW QUESTION 159

A user has enabled serverside encryption with S3. The user downloads the encrypted object from S3. How can the user decrypt it?

- A. S3 does not support server side encryption
- B. S3 provides a server side key to decrypt the object
- C. The user needs to decrypt the object using their own private key
- D. S3 manages encryption and decryption automatically

Answer: D

Explanation:

If the user is using the server-side encryption feature, Amazon S3 encrypts the object data before saving it on disks in its data centres and decrypts it when the user downloads the objects. Thus, the user is free from the tasks of managing encryption, encryption keys, and related tools.

Reference: <http://docs.aws.amazon.com/AmazonS3/latest/dev/UsingEncryption.html>

NEW QUESTION 161

An organization has enabled a strict password policy for its IAM users. The organization is taking help from the IAM console to set the password policy. Which of the below mentioned rules cannot be specified by the user as a part of the policy?

- A. Allow at least one lower case letter
- B. Allow at least one number
- C. Allow at least one non-alphanumeric character
- D. Do not allow the user to use the password from the last three passwords

Answer: D

Explanation:

AWS IAM allows an organization to create multiple users and provide them access to various AWS services. By default when the user is created, he does not have password enabled and can not login to AWS console. If the organization wants to allow the users to login to AWS console, they can enable password for each user. It is required that IAM users follow certain guidelines to set their IAM login password. For this IAM provides root account owner to setup password policy. The password policy also lets the specify whether all IAM users can change their own passwords. As part of policy, organization can specify that passwords for IAM users must be of a certain minimum length, must include certain characters, and a few more criteria such as below.

One upper/ lower or both letters One alpha numeric

One number

Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_ManagingPasswordPolicies.htm

NEW QUESTION 165

A user has developed an application which is required to send the data to a NoSQL database. The user wants to decouple the data sending such that the application keeps processing and sending data but does not wait for an acknowledgement of DB. Which of the below mentioned applications helps in this scenario?

- A. AWS Simple Notification Service
- B. AWS Simple Workflow
- C. AWS Simple Query Service
- D. AWS Simple Queue Service

Answer: D

Explanation:

Amazon Simple Queue Service (SQS) is a fast, reliable, scalable, and fully managed message queuing service. SQS provides a simple and cost-effective way to decouple the components of an application. In this case, the user can use AWS SQS to send messages which are received from an application and sent to DB. The application can continue processing data without waiting for any acknowledgement from DB. The user can use SQS to transmit any volume of data without losing messages or requiring other services to always be available.

Reference: <http://aws.amazon.com/sqs/>

NEW QUESTION 169

A user is launching an AWS RDS with MySQL. Which of the below mentioned options allows the user to configure the INNODB engine parameters?

- A. Options group
- B. Engine parameters
- C. Parameter groups
- D. DB parameters

Answer: C

Explanation:

With regard to RDS, the user can manage the configuration of a DB engine by using a DB parameter group. A DB parameter group contains engine configuration values that can be applied to one or more DB instances of the same instance type.

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

NEW QUESTION 171

ExamKiller (with AWS account ID 111122223333) has created 50 IAM users for its organization's employees. What will be the AWS console URL for these associates?

- A. <https://111122223333.signin.aws.amazon.com/console/>
- B. <https://signin.aws.amazon.com/console/>
- C. <https://signin.aws.amazon.com/111122223333/console/>
- D. <https://signin.aws.amazon.com/console/111122223333/>

Answer: A

Explanation:

When an organization is using AWS IAM for creating various users and manage their access rights, the IAM user cannot use the login URL <http://aws.amazon.com/console> to access AWS management console. The console login URL for the IAM user will have AWS account ID of that organization to identify the IAM user belongs to particular account. The AWS console login URL for the IAM user will be https://<AWS_Account_ID>.signin.aws.amazon.com/console/. In this case it will be [https:// 111122223333.signin.aws.amazon.com/console/](https://111122223333.signin.aws.amazon.com/console/)
Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/AccountAlias.html>

NEW QUESTION 174

A user is planning to host MS SQL on an EBS volume. It was recommended to use the AWS RDS. What advantages will the user have if he uses RDS in comparison to an EBS based DB?

- A. Better throughput with PIOPS
- B. Automated backup
- C. NIS SQL is not supported with RDS
- D. High availability with multi AZs

Answer: B

Explanation:

Comparing with on-premises or EC2 based NIS SQL, RDS provides an automated backup feature. PIOPS is available with both RDS and EBS. However, HA is not available with NIS SQL.
Reference: <https://aws.amazon.com/rds/faqs/>

NEW QUESTION 179

A user is planning to use the AWS RDS with MySQL. Which of the below mentioned services the user is not going to pay?

- A. Data transfer
- B. RDS Cloudwatch metrics
- C. Data storage
- D. I/O requests per month

Answer: B

Explanation:

RDS charges the user on a pay as you go basis. It charges the user based on the instance type, number of hours that the instance is running, data transfer, storage cost as well for the I/O requests. The monitoring is free of cost.
Reference: <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Welcome.html>

NEW QUESTION 181

A user has created a snapshot of an EBS volume. Which of the below mentioned usage cases is not possible with respect to a snapshot?

- A. Mirroring the volume from one AZ to another AZ
- B. Launch an instance
- C. Decrease the volume size
- D. Increase the size of the volume

Answer: C

Explanation:

The EBS snapshots are a point in time backup of the volume. It is helpful to move the volume from one AZ to another or launch a new instance. The user can increase the size of the volume but cannot decrease it less than the original snapshot size.
Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSSnapshots.html>

NEW QUESTION 184

When a user is launching an instance with EC2, which of the below mentioned options is not available during the instance launch console for a key pair?

- A. Proceed without the key pair
- B. Upload a new key pair
- C. Select an existing key pair
- D. Create a new key pair

Answer: B

Explanation:

While launching an EC2 instance, the user can create a new key pair, select an existing key pair or proceed without a key pair. The user cannot upload a new key pair in the EC2 instance launch console. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/launching-instance.html>

NEW QUESTION 187

A user is creating an EBS volume. He asks for your advice. Which advice mentioned below should you not give to the user for creating an EBS volume?

- A. Take the snapshot of the volume when the instance is stopped
- B. Stripe multiple volumes attached to the same instance
- C. Create an AMI from the attached volume
- D. Attach multiple volumes to the same instance

Answer: C

Explanation:

When a user creates an EBS volume, the user can attach it to a running instance. The user can attach multiple volumes to the same instance and stripe them together to increase the I/O. The user can take a snapshot from the existing volume but cannot create an AMI from the volume. However, the user can create an AMI from a snapshot.

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

NEW QUESTION 189

In regards to VPC, select the correct statement:

- A. You can associate multiple subnets with the same Route Table.
- B. You can associate multiple subnets with the same Route Table, but you can't associate a subnet with only one Route Table.
- C. You can't associate multiple subnets with the same Route Table.
- D. None of these

Answer: A

Explanation:

Every subnet in your VPC must be associated with exactly one Route Table. However, multiple subnets can be associated with the same Route Table.

Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Route_Tables.html

NEW QUESTION 192

A user has hosted a website on AWS and uses ELB to load balance the multiple instances. The user application does not have any cookie management. How can the user bind the session of the requestor with a particular instance?

- A. Bind the IP address with a sticky cookie
- B. Create a cookie at the application level to set at ELB
- C. Use session synchronization with ELB
- D. Let ELB generate a cookie for a specified duration

Answer: D

Explanation:

The key to manage the sticky session is determining how long the load balancer should route the user's request to the same application instance. If the application has its own session cookie, then the user can set the Elastic Load Balancing to create the session cookie to follow the duration specified by the application's session cookie. If the user's application does not have its own session cookie, then he can set the Elastic Load Balancing to create a session cookie by specifying his own stickiness duration. Reference: http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/US_StickySessions.html

NEW QUESTION 195

A user has enabled automated backup for an RDS instance. What is the longest duration for which the user can retain the automated backup?

- A. 25 days
- B. 15 days
- C. 45 days
- D. 35 days

Answer: D

Explanation:

Amazon RDS provides two different methods for backing up and restoring the Amazon DB instances: automated backups and DB snapshots. Automated backups automatically back up the DB instance during a specific, user-definable backup window, and keep the backups for a limited, user-specified period of time. The maximum period can be 35 days.

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

NEW QUESTION 198

An organization has 10 departments. The organization wants to track the AWS usage of each department. Which of the below mentioned options meets the requirement?

- A. Setup IAM groups for each department and track their usage
- B. Create separate accounts for each department, but use consolidated billing for payment and tracking
- C. Create separate accounts for each department and track them separately
- D. Setup IAM users for each department and track their usage

Answer: B

Explanation:

The cost of an IAM user or groups can never be tracked separately for the purpose of billing. The best solution in this case is to create a separate account for each department and use consolidated billing. Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/IAM_Introduction.html

NEW QUESTION 199

In regard to DynamoDB, for which one of the following parameters does Amazon not charge you?

- A. Cost per provisioned write units
- B. Cost per provisioned read units
- C. Storage cost
- D. I/O usage within the same Region

Answer: D

Explanation:

In DynamoDB, you will be charged for the storage and the throughput you use rather than for the I/O which has been used.

Reference: <http://aws.amazon.com/dynamodb/pricing/>

NEW QUESTION 204

A user is planning to host data with RDS. Which of the below mentioned databases is not supported by RDS?

- A. PostgreSQL
- B. SQLDB
- C. Oracle
- D. MS SQL

Answer: B

Explanation:

Amazon Relational Database Service (Amazon RDS) is a web service that makes it easier to set up, operate, and scale a relational database in the cloud. AWS RDS supports popular DBs, such as MySQL, PostgreSQL, MS SQL and Oracle. This means that the code, applications, and tools user is already using with existing databases can be used with Amazon RDS too. In short, it is a managed Relation Database offering from AWS which manages backups, software patching, automatic failure detection, and recovery of Database.

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

NEW QUESTION 207

An EC2 instance has one additional EBS volume attached to it. How can a user attach the same volume to another running instance in the same AZ?

- A. Terminate the first instance and only then attach to the new instance
- B. Attach the volume as read only to the second instance
- C. Detach the volume first and attach to new instance
- D. No need to detach
- E. Just select the volume and attach it to the new instance, it will take care of mapping internally

Answer: C

Explanation:

If an EBS volume is attached to a running EC2 instance, the user needs to detach the volume from the original instance and then attach it to a new running instance. The user doesn't need to stop / terminate the original instance.

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

NEW QUESTION 212

Regarding Amazon SNS, to send messages to a queue through a topic, you must subscribe the queue to the Amazon SNS topic. You specify the queue by its .

- A. ARN
- B. Token
- C. Registration ID
- D. URL

Answer: A

Explanation:

In Amazon SNS, to send messages to a queue through a topic, you must subscribe the queue to the Amazon SNS topic. You specify the queue by its ARN.

Reference: <http://docs.aws.amazon.com/sns/latest/dg/SendMessageToSQS.html>

NEW QUESTION 215

A user has configured a website and launched it using the Apache web server on port 80. The user is using ELB with the EC2 instances for Load Balancing. What should the user do to ensure that the EC2 instances accept requests only from ELB?

- A. Open the port for an ELB static IP in the EC2 security group
- B. Configure the security group of EC2, which allows access to the ELB source security group
- C. Configure the EC2 instance so that it only listens on the ELB port
- D. Configure the security group of EC2, which allows access only to the ELB listener

Answer: B

Explanation:

When a user is configuring ELB and registering the EC2 instances with it, ELB will create a source security group. If the user wants to allow traffic only from ELB, he should remove all the rules set for the other requests and open the port only for the ELB source security group.

Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/using-elb-security-groups.html>

NEW QUESTION 218

A bucket owner has allowed another account's IAM users to upload or access objects in his bucket. The IAM user of Account A is trying to access an object created by the IAM user of account B. What will happen in this scenario?

- A. The bucket policy may not be created as S3 will give error due to conflict of Access Rights
- B. It is not possible to give permission to multiple IAM users
- C. AWS S3 will verify proper rights given by the owner of Account A, the bucket owner as well as by the IAM user B to the object

D. It is not possible that the IAM user of one account accesses objects of the other IAM user

Answer: C

Explanation:

If a IAM user is trying to perform some action on an object belonging to another AWS user's bucket, S3 will verify whether the owner of the IAM user has given sufficient permission to him. It also verifies the policy for the bucket as well as the policy defined by the object owner.

Reference:

<http://docs.aws.amazon.com/AmazonS3/latest/dev/access-control-auth-workflow-object-operation.html>

NEW QUESTION 219

A user wants to achieve High Availability with PostgreSQL DB. Which of the below mentioned functionalities helps achieve HA?

- A. Read Replica
- B. Multi AZ
- C. Multi region
- D. PostgreSQL does not support HA

Answer: B

Explanation:

The Multi AZ feature allows the user to achieve High Availability. For Multi AZ, Amazon RDS automatically provisions and maintains a synchronous "standby" replica in a different Availability Zone. Reference:

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

NEW QUESTION 222

Which statements about DynamoDB are true? Choose 2 answers

- A. DynamoDB uses a pessimistic locking model
- B. DynamoDB uses optimistic concurrency control
- C. DynamoDB uses conditional writes for consistency
- D. DynamoDB restricts item access during reads
- E. DynamoDB restricts item access during writes

Answer: BC

NEW QUESTION 227

You have an environment that consists of a public subnet using Amazon VPC and 3 instances that are running in this subnet. These three instances can successfully communicate with other hosts on the Internet. You launch a fourth instance in the same subnet, using the same AMI and security group configuration you used for the others, but find that this instance cannot be accessed from the Internet. What should you do to enable internet access?

- A. Deploy a NAT instance into the public subnet.
- B. Modify the routing table for the public subnet
- C. Configure a publically routable IP Address In the host OS of the fourth instance.
- D. Assign an Elastic IP address to the fourth instanc

Answer: D

NEW QUESTION 231

Company C has recently launched an online commerce site for bicycles on AWS. They have a "Product" DynamoDB table that stores details for each bicycle, such as, manufacturer, color, price, quantity and size to display in the online store. Due to customer demand, they want to include an image for each bicycle along with the existing details.

Which approach below provides the least impact to provisioned throughput on the "Product" table?

- A. Serialize the image and store it in multiple DynamoDB tables
- B. Create an "Images" DynamoDB table to store the Image with a foreign key constraint to the "Product" table
- C. Add an image data type to the "Product" table to store the images in binary format
- D. Store the images in Amazon S3 and add an S3 URL pointer to the "Product" table item for each image

Answer: D

NEW QUESTION 236

Company D is running their corporate website on Amazon S3 accessed from <http://www.companyd.com>. Their marketing team has published new web fonts to a separate S3 bucket accessed by the S3 endpoint <https://s3-us-west1.amazonaws.com/cdfonts>. While testing the new web fonts, Company D recognized the web fonts are being blocked by the browser. What should Company D do to prevent the web fonts from being blocked by the browser?

- A. Enable versioning on the cdfonts bucket for each web font
- B. Create a policy on the cdfonts bucket to enable access to everyone
- C. Add the Content-NI D5 header to the request for webfonts in the cdfonts bucket from the website
- D. Configure the cdfonts bucket to allow cross-origin requests by creating a CORS configuration

Answer: D

NEW QUESTION 239

A meteorological system monitors 600 temperature gauges, obtaining temperature samples every minute and saving each sample to a DynamoDB table. Each sample involves writing 1K of data and the writes are evenly distributed over time.

How much write throughput is required for the target table?

- A. 1 write capacity unit
- B. 10 write capacity units
- C. 60 write capacity units
- D. 600 write capacity units
- E. 3600 write capacity units

Answer: B

NEW QUESTION 242

You have written an application that uses the Elastic Load Balancing service to spread traffic to several web servers. Your users complain that they are sometimes forced to login again in the middle of using your application, after they have already logged in. This is not behavior you have designed. What is a possible solution to prevent this happening?

- A. Use instance memory to save session state.
- B. Use instance storage to save session state.
- C. Use EBS to save session state
- D. Use ElastiCache to save session state.
- E. Use Glacier to save session state

Answer: D

NEW QUESTION 247

Which DynamoDB limits can be raised by contacting AWS support? Choose 2 answers

- A. The number of hash keys per account
- B. The maximum storage used per account
- C. The number of tables per account
- D. The number of local secondary indexes per account
- E. The number of provisioned throughput units per account

Answer: CE

NEW QUESTION 251

In DynamoDB, what type of HTTP response codes indicate that a problem was found with the client request sent to the service?

- A. 5xx HTTP response code
- B. 200 HTTP response code
- C. 306 HTTP response code
- D. 4xx HTTP response code

Answer: D

NEW QUESTION 253

What item operation allows the retrieval of multiple items from a DynamoDB table in a single API call?

- A. GetItem
- B. BatchGetItem
- C. GetMultipleItems
- D. GetItemRange

Answer: B

NEW QUESTION 254

An application stores payroll information nightly in DynamoDB for a large number of employees across hundreds of offices. Item attributes consist of individual name, office identifier, and cumulative daily hours. Managers run reports for ranges of names working in their office. One query is, "Return all items in this office for names starting with A through E". Which table configuration will result in the lowest impact on provisioned throughput for this query?

- A. Configure the table to have a hash index on the name attribute, and a range index on the office identifier
- B. Configure the table to have a range index on the name attribute, and a hash index on the office identifier
- C. Configure a hash index on the name attribute and no range index
- D. Configure a hash index on the office identifier attribute and no range index

Answer: B

NEW QUESTION 256

EC2 instances are launched from Amazon Machine images (AMIs). A given public AMI can:

- A. be used to launch EC2 instances in any AWS region.
- B. only be used to launch EC2 instances in the same country as the AMI is stored.
- C. only be used to launch EC2 instances in the same AWS region as the AMI is stored.
- D. only be used to launch EC2 instances in the same AWS availability zone as the AMI is stored

Answer: C

NEW QUESTION 258

Which EC2 API call would you use to retrieve a list of Amazon Machine Images (AMIs)?

- A. DescrbelInstances
- B. DescribeAMIs
- C. DescribeImages
- D. GetAMIs
- E. You cannot retrieve a list of AMIs as there are over 10,000 AMIs

Answer: E

NEW QUESTION 262

What is the format of structured notification messages sent by Amazon SNS?

- A. An XML object containing MessageId, UnsubscribeURL, Subject, IVmessage and other values
- B. An JSON object containing MessageId, DuplicateFlag, IVmessage and other values
- C. An XML object containing MessageId, DuplicateFlag, IVmessage and other values
- D. An JSON object containing MessageId, unsubscribeURL, Subject, IVmessage and other values

Answer: D

NEW QUESTION 264

When using a large Scan operation in DynamoDB, what technique can be used to minimize the impact of a scan on a table's provisioned throughput?

- A. Set a smaller page size for the scan
- B. Use parallel scans
- C. Define a range index on the table
- D. Prewarm the table by updating all items

Answer: C

NEW QUESTION 265

How can software determine the public and private IP addresses of the Amazon EC2 instance that it is running on?

- A. Query the appropriate Amazon CloudWatch metric.
- B. Use ipconfig or ifconfig command.
- C. Query the local instance userdata.
- D. Query the local instance metadata

Answer: D

NEW QUESTION 268

An Amazon S3 bucket, "myawsbucket" is configured with website hosting in Tokyo region, what is the region-specific website endpoint?

- A. www.myawsbucket.ap-northeast-1.amazonaws.com
- B. myawsbucket.s3-website-ap-northeast-1.amazonaws.com
- C. myawsbucket.amazonaws.com
- D. myawsbucket.tokyo.amazonaws.com

Answer: B

NEW QUESTION 271

Your application is trying to upload a 6 GB file to Simple Storage Service and receive a "Your proposed upload exceeds the maximum allowed object size." error message.

What is a possible solution for this?

- A. None, Simple Storage Service objects are limited to 5 GB
- B. Use the multi-part upload API for this object
- C. Use the large object upload API for this object
- D. Contact support to increase your object size limit
- E. Upload to a different region

Answer: B

NEW QUESTION 272

What type of block cipher does Amazon S3 offer for server side encryption?

- A. Triple DES
- B. Advanced Encryption Standard
- C. Blowfish
- D. RC5

Answer: B

NEW QUESTION 274

You attempt to store an object in the US-STANDARD region in Amazon S3, and receive a confirmation that it has been successfully stored. You then immediately make another API call and attempt to read this object. S3 tells you that the object does not exist. What could explain this behavior?

- A. US-STANDARD uses eventual consistency and it can take time for an object to be readable in a bucket
- B. Objects in Amazon S3 do not become visible until they are replicated to a second region.
- C. US-STANDARD imposes a 1 second delay before new objects are readable.
- D. You exceeded the bucket object limit, and once this limit is raised the object will be visible.

Answer: A

NEW QUESTION 275

After launching an instance that you intend to serve as a NAT (Network Address Translation) device in a public subnet you modify your route tables to have the NAT device be the target of internet bound traffic of your private subnet. When you try and make an outbound connection to the Internet from an instance in the private subnet, you are not successful.

Which of the following steps could resolve the issue?

- A. Attaching a second Elastic Network interface (ENI) to the NAT instance, and placing it in the private subnet
- B. Attaching a second Elastic Network Interface (ENI) to the instance in the private subnet, and placing it in the public subnet
- C. Disabling the Source/Destination Check attribute on the NAT instance
- D. Attaching an Elastic IP address to the instance in the private subnet

Answer: C

NEW QUESTION 276

Which of the following statements about SQS is true?

- A. Messages will be delivered exactly once and messages will be delivered in First in, First out order
- B. Messages will be delivered exactly once and message delivery order is indeterminate
- C. Messages will be delivered one or more times and messages will be delivered in First in, First out order
- D. Messages will be delivered one or more times and message delivery order is indeterminate

Answer: D

NEW QUESTION 279

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual DVA-C01 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the DVA-C01 Product From:

<https://www.2passeasy.com/dumps/DVA-C01/>

Money Back Guarantee

DVA-C01 Practice Exam Features:

- * DVA-C01 Questions and Answers Updated Frequently
- * DVA-C01 Practice Questions Verified by Expert Senior Certified Staff
- * DVA-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * DVA-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year