

AWS-Solution-Architect-Associate Dumps

Amazon AWS Certified Solutions Architect - Associate

<https://www.certleader.com/AWS-Solution-Architect-Associate-dumps.html>



NEW QUESTION 1

You are trying to launch an EC2 instance, however the instance seems to go into a terminated status immediately. What would probably not be a reason that this is happening?

- A. The AMI is missing a required part.
- B. The snapshot is corrupt.
- C. You need to create storage in EBS first.
- D. You've reached your volume limit

Answer: C

Explanation:

Amazon EC2 provides a virtual computing environments, known as an instance.

After you launch an instance, AWS recommends that you check its status to confirm that it goes from the pending status to the running status, the not terminated status.

The following are a few reasons why an Amazon EBS-backed instance might immediately terminate: You've reached your volume limit.

The AMI is missing a required part. The snapshot is corrupt. Reference:

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

NEW QUESTION 2

In the context of AWS support, why must an EC2 instance be unreachable for 20 minutes rather than allowing customers to open tickets immediately?

- A. Because most reachability issues are resolved by automated processes in less than 20 minutes
- B. Because all EC2 instances are unreachable for 20 minutes every day when AWS does routine maintenance
- C. Because all EC2 instances are unreachable for 20 minutes when first launched
- D. Because of all the reasons listed here

Answer: A

Explanation:

An EC2 instance must be unreachable for 20 minutes before opening a ticket, because most reachability issues are resolved by automated processes in less than 20 minutes and will not require any action on the part of the customer. If the instance is still unreachable after this time frame has passed, then you should open a case with support.

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

NEW QUESTION 3

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

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

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 4

To specify a resource in a policy statement, in Amazon EC2, can you use its Amazon Resource Name (ARN)?

- A. Yes, you can.
- B. No, you can't because EC2 is not related to ARN.
- C. No, you can't because you can't specify a particular Amazon EC2 resource in an IAM policy.
- D. Yes, you can but only for the resources that are not affected by the action

Answer: A

Explanation:

Some Amazon EC2 API actions allow you to include specific resources in your policy that can be created or modified by the action. To specify a resource in the statement, you need to use its Amazon Resource Name (ARN).

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-ug.pdf>

NEW QUESTION 5

After you recommend Amazon Redshift to a client as an alternative solution to paying data warehouses to analyze his data, your client asks you to explain why you are recommending Redshift. Which of the following would be a reasonable response to his request?

- A. It has high performance at scale as data and query complexity grows.
- B. It prevents reporting and analytic processing from interfering with the performance of OLTP workloads.
- C. You don't have the administrative burden of running your own data warehouse and dealing with setup, durability, monitoring, scaling, and patching.
- D. All answers listed are a reasonable response to his QUESTION

Answer: D

Explanation:

Amazon Redshift delivers fast query performance by using columnar storage technology to improve I/O efficiency and parallelizing queries across multiple nodes. Redshift uses standard PostgreSQL JDBC and ODBC drivers, allowing you to use a wide range of familiar SQL clients. Data load speed scales linearly with cluster size, with integrations to Amazon S3, Amazon DynamoDB, Amazon Elastic MapReduce, Amazon Kinesis or any SSH-enabled host.

AWS recommends Amazon Redshift for customers who have a combination of needs, such as: High performance at scale as data and query complexity grows
Desire to prevent reporting and analytic processing from interfering with the performance of OLTP workloads

Large volumes of structured data to persist and query using standard SQL and existing BI tools
Desire to the administrative burden of running one's own data warehouse and dealing with setup, durability, monitoring, scaling and patching

Reference: https://aws.amazon.com/running_databases/#redshift_anchor

NEW QUESTION 6

A user is storing a large number of objects on AWS S3. The user wants to implement the search functionality among the objects. How can the user achieve this?

- A. Use the indexing feature of S3.
- B. Tag the objects with the metadata to search on that.
- C. Use the query functionality of S3.
- D. Make your own DB system which stores the S3 metadata for the search functionalit

Answer: D

Explanation:

In Amazon Web Services, AWS S3 does not provide any query facility. To retrieve a specific object the user needs to know the exact bucket / object key. In this case it is recommended to have an own DB system which manages the S3 metadata and key mapping.

Reference: http://media.amazonwebservices.com/AWS_Storage_Options.pdf

NEW QUESTION 7

A client needs you to import some existing infrastructure from a dedicated hosting provider to AWS to try and save on the cost of running his current website. He also needs an automated process that manages backups, software patching, automatic failure detection, and recovery. You are aware that his existing set up currently uses an Oracle database. Which of the following AWS databases would be best for accomplishing this task?

- A. Amazon RDS
- B. Amazon Redshift
- C. Amazon SimpleDB
- D. Amazon ElastiCache

Answer: A

Explanation:

Amazon RDS gives you access to the capabilities of a familiar MySQL, Oracle, SQL Server, or PostgreSQL database engine. This means that the code, applications, and tools you already use today with your existing databases can be used with Amazon RDS. Amazon RDS automatically patches the database software and backs up your database, storing the backups for a user-defined retention period and enabling point-in-time recovery.

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

NEW QUESTION 8

An edge location refers to which Amazon Web Service?

- A. An edge location is referred to the network configured within a Zone or Region
- B. An edge location is an AWS Region
- C. An edge location is the location of the data center used for Amazon CloudFront.
- D. An edge location is a Zone within an AWS Region

Answer: C

Explanation:

Amazon CloudFront is a content distribution network. A content delivery network or content distribution network (CDN) is a large distributed system of servers deployed in multiple data centers across the world. The location of the data center used for CDN is called edge location.

Amazon CloudFront can cache static content at each edge location. This means that your popular static content (e.g., your site's logo, navigational images, cascading style sheets, JavaScript code, etc.) will be available at a nearby edge location for the browsers to download with low latency and improved performance for viewers. Caching popular static content with Amazon CloudFront also helps you offload requests for such files from your origin server — CloudFront serves the cached copy when available and only makes a request to your origin server if the edge location receiving the browser's request does not have a copy of the file.

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

NEW QUESTION 9

Does Amazon DynamoDB support both increment and decrement atomic operations?

- A. Only increment, since decrement are inherently impossible with DynamoDB's data model.
- B. No, neither increment nor decrement operations.
- C. Yes, both increment and decrement operations.
- D. Only decrement, since increment are inherently impossible with DynamoDB's data mode

Answer: C

Explanation:

Amazon DynamoDB supports increment and decrement atomic operations.

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

NEW QUESTION 10

A user has created an EBS volume with 1000 IOPS. What is the average IOPS that the user will get for most of the year as per EC2 SLA if the instance is attached

to the EBS optimized instance?

- A. 950
- B. 990
- C. 1000
- D. 900

Answer: D

Explanation:

As per AWS SLA if the instance is attached to an EBS-Optimized instance, then the Provisioned IOPS volumes are designed to deliver within 10% of the provisioned IOPS performance 99.9% of the time in a given year. Thus, if the user has created a volume of 1000 IOPS, the user will get a minimum 900 IOPS 99.9% time of the year.

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

NEW QUESTION 10

You need to migrate a large amount of data into the cloud that you have stored on a hard disk and you decide that the best way to accomplish this is with AWS Import/Export and you mail the hard disk to AWS. Which of the following statements is incorrect in regards to AWS Import/Export?

- A. It can export from Amazon S3
- B. It can Import to Amazon Glacier
- C. It can export from Amazon Glacier.
- D. It can Import to Amazon EBS

Answer: C

Explanation:

AWS Import/Export supports: Import to Amazon S3

Export from Amazon S3 Import to Amazon EBS Import to Amazon Glacier

AWS Import/Export does not currently support export from Amazon EBS or Amazon Glacier. Reference:

<https://docs.aws.amazon.com/AWSImportExport/latest/DG/whatisdisk.html>

NEW QUESTION 14

In EC2, what happens to the data in an instance store if an instance reboots (either intentionally or unintentionally)?

- A. Data is deleted from the instance store for security reasons.
- B. Data persists in the instance store.
- C. Data is partially present in the instance store.
- D. Data in the instance store will be los

Answer: B

Explanation:

The data in an instance store persists only during the lifetime of its associated instance. If an instance reboots (intentionally or unintentionally), data in the instance store persists. However, data on instance store volumes is lost under the following circumstances.

Failure of an underlying drive

Stopping an Amazon EBS-backed instance Terminating an instance

Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/InstanceStorage.html>

NEW QUESTION 19

Can you specify the security group that you created for a VPC when you launch an instance in EC2-Classic?

- A. No, you can specify the security group created for EC2-Classic when you launch a VPC instance.
- B. No
- C. Yes
- D. No, you can specify the security group created for EC2-Classic to a non-VPC based instance onl

Answer: B

Explanation:

If you're using EC2-Classic, you must use security groups created specifically for EC2-Classic. When you launch an instance in EC2-Classic, you must specify a security group in the same region as the instance. You can't specify a security group that you created for a VPC when you launch an instance in EC2-Classic.

Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html#ec2-classic-security-groups>

NEW QUESTION 20

While using the EC2 GET requests as URLs, the is the URL that serves as the entry point for the web service.

- A. token
- B. endpoint
- C. action
- D. None of these

Answer: B

Explanation:

The endpoint is the URL that serves as the entry point for the web service.

Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/using-query-api.html>

NEW QUESTION 21

Which of the below mentioned options is not available when an instance is launched by Auto Scaling with EC2 Classic?

- A. Public IP
- B. Elastic IP
- C. Private DNS
- D. Private IP

Answer: B

Explanation:

Auto Scaling supports both EC2 classic and EC2-VPC. When an instance is launched as a part of EC2 classic, it will have the public IP and DNS as well as the private IP and DNS.

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

NEW QUESTION 26

Much of your company's data does not need to be accessed often, and can take several hours for retrieval time, so it's stored on Amazon Glacier. However someone within your organization has expressed concerns that his data is more sensitive than the other data, and is wondering whether the high level of encryption that he knows is on S3 is also used on the much cheaper Glacier service. Which of the following statements would be most applicable in regards to this concern?

- A. There is no encryption on Amazon Glacier, that's why it is cheaper.
- B. Amazon Glacier automatically encrypts the data using AES-128 a lesser encryption method than Amazon S3 but you can change it to AES-256 if you are willing to pay more.
- C. Amazon Glacier automatically encrypts the data using AES-256, the same as Amazon S3.
- D. Amazon Glacier automatically encrypts the data using AES-128 a lesser encryption method than Amazon S3.

Answer: C

Explanation:

Like Amazon S3, the Amazon Glacier service provides low-cost, secure, and durable storage. But where S3 is designed for rapid retrieval, Glacier is meant to be used as an archival service for data that is not accessed often, and for which retrieval times of several hours are suitable.

Amazon Glacier automatically encrypts the data using AES-256 and stores it durably in an immutable form. Amazon Glacier is designed to provide average annual durability of 99.999999999% for an archive. It stores each archive in multiple facilities and multiple devices. Unlike traditional systems which can require laborious data verification and manual repair, Glacier performs regular, systematic data integrity checks, and is built to be automatically self-healing.

Reference: <http://d0.awsstatic.com/whitepapers/Security/AWS%20Security%20Whitepaper.pdf>

NEW QUESTION 27

Your EBS volumes do not seem to be performing as expected and your team leader has requested you look into improving their performance. Which of the following is not a true statement relating to the performance of your EBS volumes?

- A. Frequent snapshots provide a higher level of data durability and they will not degrade the performance of your application while the snapshot is in progress.
- B. General Purpose (SSD) and Provisioned IOPS (SSD) volumes have a throughput limit of 128 MB/s per volume.
- C. There is a relationship between the maximum performance of your EBS volumes, the amount of I/O you are drMng to them, and the amount of time it takes for each transaction to complete.
- D. There is a 5 to 50 percent reduction in IOPS when you first access each block of data on a newly created or restored EBS volume

Answer: A

Explanation:

Several factors can affect the performance of Amazon EBS volumes, such as instance configuration, I/O characteristics, workload demand, and storage configuration.

Frequent snapshots provide a higher level of data durability, but they may slightly degrade the performance of your application while the snapshot is in progress. This trade off becomes critical when you have data that changes rapidly. Whenever possible, plan for snapshots to occur during off-peak times in order to minimize workload impact.

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

NEW QUESTION 29

In Amazon EC2 Container Service, are other container types supported?

- A. Yes, EC2 Container Service supports any container service you need.
- B. Yes, EC2 Container Service also supports Microsoft container service.
- C. No, Docker is the only container platform supported by EC2 Container Service presently.
- D. Yes, EC2 Container Service supports Microsoft container service and Openstac

Answer: C

Explanation:

In Amazon EC2 Container Service, Docker is the only container platform supported by EC2 Container Service presently.

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

NEW QUESTION 31

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

Answer: C

Explanation:

Billing commences when Amazon EC2 initiates the boot sequence of an AM 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 35

You want to use AWS Import/Export to send data from your S3 bucket to several of your branch offices. What should you do if you want to send 10 storage units to AWS?

- A. Make sure your disks are encrypted prior to shipping.
- B. Make sure you format your disks prior to shipping.
- C. Make sure your disks are 1TB or more.
- D. Make sure you submit a separate job request for each device

Answer: D

Explanation:

When using Amazon Import/Export, a separate job request needs to be submitted for each physical device even if they belong to the same import or export job.

Reference: <http://docs.aws.amazon.com/AWSImportExport/latest/DG/Concepts.html>

NEW QUESTION 39

What would be the best way to retrieve the public IP address of your EC2 instance using the CLI?

- A. Using tags
- B. Using traceroute
- C. Using ipconfig
- D. Using instance metadata

Answer: D

Explanation:

To determine your instance's public IP address from within the instance, you can use instance metadata. Use the following command to access the public IP address: For Linux use, `$ curl`

`http://169.254.169.254/latest/meta-data/public-ipv4`, and for Windows use, `$ wget http://169.254.169.254/latest/meta-data/public-ipv4`.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-instance-addressing.htm>

NEW QUESTION 43

You need to measure the performance of your EBS volumes as they seem to be under performing. You have come up with a measurement of 1,024 KB I/O but your colleague tells you that EBS volume performance is measured in IOPS. How many IOPS is equal to 1,024 KB I/O?

- A. 16
- B. 256
- C. 8
- D. 4

Answer: D

Explanation:

Several factors can affect the performance of Amazon EBS volumes, such as instance configuration, I/O characteristics, workload demand, and storage configuration.

IOPS are input/output operations per second. Amazon EBS measures each I/O operation per second

(that is 256 KB or smaller) as one IOPS. I/O operations that are larger than 256 KB are counted in 256 KB capacity units.

For example, a 1,024 KB I/O operation would count as 4 IOPS.

When you provision a 4,000 IOPS volume and attach it to an EBS-optimized instance that can provide the necessary bandwidth, you can transfer up to 4,000 chunks of data per second (provided that the I/O does not exceed the 128 MB/s per volume throughput limit of General Purpose (SSD) and Provisioned IOPS (SSD) volumes).

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

NEW QUESTION 48

You decide that you need to create a number of Auto Scaling groups to try and save some money as you have noticed that at certain times most of your EC2 instances are not being used. By default, what is the maximum number of Auto Scaling groups that AWS will allow you to create?

- A. 12
- B. Unlimited
- C. 20
- D. 2

Answer: C

Explanation:

Auto Scaling is an AWS service that allows you to increase or decrease the number of EC2 instances within your application's architecture. With Auto Scaling, you

create collections of EC2 instances, called Auto Scaling groups. You can create these groups from scratch, or from existing EC2 instances that are already in production.

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

NEW QUESTION 50

You are setting up your first Amazon Virtual Private Cloud (Amazon VPC) so you decide to use the VPC wizard in the AWS console to help make it easier for you. Which of the following statements is correct regarding instances that you launch into a default subnet via the VPC wizard?

- A. Instances that you launch into a default subnet receive a public IP address and 10 private IP addresses.
- B. Instances that you launch into a default subnet receive both a public IP address and a private IP address.
- C. Instances that you launch into a default subnet don't receive any IP addresses and you need to define them manually.
- D. Instances that you launch into a default subnet receive a public IP address and 5 private IP addresses.

Answer: B

Explanation:

Instances that you launch into a default subnet receive both a public IP address and a private IP address. Instances in a default subnet also receive both public and private DNS hostnames. Instances that you launch into a nondefault subnet in a default VPC don't receive a public IP address or a DNS hostname. You can change your subnet's default public IP addressing behavior.

Reference: <http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/default-vpc.html>

NEW QUESTION 54

A user has configured ELB with two EBS backed EC2 instances. The user is trying to understand the DNS access and IP support for ELB. Which of the below mentioned statements may not help the user understand the IP mechanism supported by ELB?

- A. The client can connect over IPV4 or IPV6 using Dualstack
- B. Communication between the load balancer and back-end instances is always through IPV4
- C. ELB DNS supports both IPV4 and IPV6
- D. The ELB supports either IPV4 or IPV6 but not both

Answer: D

Explanation:

Elastic Load Balancing supports both Internet Protocol version 6 (IPv6) and Internet Protocol version 4 (IPv4). Clients can connect to the user's load balancer using either IPv4 or IPv6 (in EC2-Classic) DNS. However, communication between the load balancer and its back-end instances uses only IPv4. The user can use the Dualstack-prefixed DNS name to enable IPv6 support for communications between the client and the load balancers. Thus, the clients are able to access the load balancer using either IPv4 or IPv6 as their individual connectivity needs dictate.

Reference: <http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/UserScenariosForEC2.html>

NEW QUESTION 59

Amazon S3 allows you to set per-file permissions to grant read and/or write access. However you have decided that you want an entire bucket with 100 files already in it to be accessible to the public. You don't want to go through 100 files individually and set permissions. What would be the best way to do this?

- A. Move the bucket to a new region
- B. Add a bucket policy to the bucket.
- C. Move the files to a new bucket.
- D. Use Amazon EBS instead of S3

Answer: B

Explanation:

Amazon S3 supports several mechanisms that give you flexibility to control who can access your data as well as how, when, and where they can access it. Amazon S3 provides four different access control mechanisms: AWS Identity and Access Management (IAM) policies, Access Control Lists (ACLs), bucket policies, and query string authentication. IAM enables organizations to create and manage multiple users under a single AWS account. With IAM policies, you can grant IAM users fine-grained control to your Amazon S3 bucket or objects. You can use ACLs to selectively add (grant) certain permissions on individual objects. Amazon S3 bucket policies can be used to add or deny permissions across some or all of the objects within a single bucket. With Query string authentication, you have the ability to share Amazon S3 objects through URLs that are valid for a specified period of time.

Reference: <http://aws.amazon.com/s3/details/#security>

NEW QUESTION 60

You need to set up a high level of security for an Amazon Relational Database Service (RDS) you have just built in order to protect the confidential information stored in it. What are all the possible security groups that RDS uses?

- A. DB security groups, VPC security groups, and EC2 security groups.
- B. DB security groups only.
- C. EC2 security groups only.
- D. VPC security groups, and EC2 security group

Answer: A

Explanation:

A security group controls the access to a DB instance. It does so by allowing access to IP address ranges or Amazon EC2 instances that you specify. Amazon RDS uses DB security groups, VPC security groups, and EC2 security groups. In simple terms, a DB security group controls access to a DB instance that is not in a VPC, a VPC security group controls access to a DB instance inside a VPC, and an Amazon EC2 security group controls access to an EC2 instance and can be used with a DB instance.

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

NEW QUESTION 61

You have been using T2 instances as your CPU requirements have not been that intensive. However you now start to think about larger instance types and start looking at M and IV|3 instances. You are a little confused as to the differences between them as they both seem to have the same ratio of CPU and memory. Which statement below is incorrect as to why you would use one over the other?

- A. M3 instances are less expensive than M1 instances.
- B. IV|3 instances are configured with more swap memory than M instances.
- C. IV|3 instances provide better, more consistent performance than M instances for most use-cases.
- D. M3 instances also offer SSD-based instance storage that delivers higher I/O performance.

Answer: B

Explanation:

Amazon EC2 allows you to set up and configure everything about your instances from your operating system up to your applications. An Amazon Machine Image (AMI) is simply a packaged-up environment that includes all the necessary bits to set up and boot your instance. M1 and M3 Standard instances have the same ratio of CPU and memory, some reasons below as to why you would use one over the other. IV|3 instances provide better, more consistent performance than M instances for most use-cases. M3 instances also offer SSD-based instance storage that delivers higher I/O performance. M3 instances are also less expensive than M1 instances. Due to these reasons, we recommend M3 for applications that require general purpose instances with a balance of compute, memory, and network resources. However, if you need more disk storage than what is provided in M3 instances, you may still find M1 instances useful for running your applications. Reference: <https://aws.amazon.com/ec2/faqs/>

NEW QUESTION 63

A user has created an application which will be hosted on EC2. The application makes calls to DynamoDB to fetch certain data. The application is using the DynamoDB SDK to connect with from the EC2 instance. Which of the below mentioned statements is true with respect to the best practice for security in this scenario?

- A. The user should create an IAM user with DynamoDB access and use its credentials within the application to connect with DynamoDB
- B. The user should attach an IAM role with DynamoDB access to the EC2 instance
- C. The user should create an IAM role, which has EC2 access so that it will allow deploying the application
- D. The user should create an IAM user with DynamoDB and EC2 access
- E. Attach the user with the application so that it does not use the root account credentials

Answer: B

Explanation:

With AWS IAM a user is creating an application which runs on an EC2 instance and makes requests to AWS, such as DynamoDB or S3 calls. Here it is recommended that the user should not create an IAM user and pass the user's credentials to the application or embed those credentials inside the application. Instead, the user should use roles for EC2 and give that role access to DynamoDB / S3. When the roles are attached to EC2, it will give temporary security credentials to the application hosted on that EC2, to connect with DynamoDB / S3. Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_WorkingWithGroupsAndUsers.html

NEW QUESTION 66

After setting up several database instances in Amazon Relational Database Service (Amazon RDS) you decide that you need to track the performance and health of your databases. How can you do this?

- A. Subscribe to Amazon RDS events to be notified when changes occur with a DB instance, DB snapshot, DB parameter group, or DB security group.
- B. Use the free Amazon CloudWatch service to monitor the performance and health of a DB instance.
- C. All of the items listed will track the performance and health of a database.
- D. View, download, or watch database log files using the Amazon RDS console or Amazon RDS API
- E. You can also query some database log files that are loaded into database tables.

Answer: C

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. It provides cost-efficient, resizable capacity for an industry-standard relational database and manages common database administration tasks. There are several ways you can track the performance and health of a database or a DB instance. You can: Use the free Amazon CloudWatch service to monitor the performance and health of a DB instance. Subscribe to Amazon RDS events to be notified when changes occur with a DB instance, DB snapshot, DB parameter group, or DB security group. View, download, or watch database log files using the Amazon RDS console or Amazon RDS APIs. You can also query some database log files that are loaded into database tables. Use the AWS CloudTrail service to record AWS calls made by your AWS account. The calls are recorded in log files and stored in an Amazon S3 bucket. Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Monitoring.html

NEW QUESTION 69

A user has created a subnet in VPC and launched an EC2 instance within it. The user has not selected the option to assign the IP address while launching the instance. The user has 3 elastic IPs and is trying to assign one of the Elastic IPs to the VPC instance from the console. The console does not show any instance in the IP assignment screen. What is a possible reason that the instance is unavailable in the assigned IP console?

- A. The IP address may be attached to one of the instances
- B. The IP address belongs to a different zone than the subnet zone
- C. The user has not created an internet gateway
- D. The IP addresses belong to EC2 Classic; so they cannot be assigned to VPC

Answer: D

Explanation:

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that

subnet. When the user is launching an instance he needs to select an option which attaches a public IP to the instance. If the user has not selected the option to attach the public IP then it will only have a private IP when launched. If the user wants to connect to an instance from the internet he should create an elastic IP with VPC. If the elastic IP is a part of EC2 Classic it cannot be assigned to a VPC instance.

Reference: <http://docs.aws.amazon.com/AmazonVPC/latest/GettingStartedGuide/LaunchInstance.html>

NEW QUESTION 73

An accountant asks you to design a small VPC network for him and, due to the nature of his business, just needs something where the workload on the network will be low, and dynamic data will be accessed infrequently. Being an accountant, low cost is also a major factor. Which EBS volume type would best suit his requirements?

- A. Magnetic
- B. Any, as they all perform the same and cost the same.
- C. General Purpose (SSD)
- D. Magnetic or Provisioned IOPS (SSD)

Answer: A

Explanation:

You can choose between three EBS volume types to best meet the needs of their workloads: General Purpose (SSD), Provisioned IOPS (SSD), and Magnetic. General Purpose (SSD) is the new, SSD-backed, general purpose EBS volume type that we recommend as the default choice for customers. General Purpose (SSD) volumes are suitable for a broad range of workloads, including small to medium sized databases, development and test environments, and boot volumes. Provisioned IOPS (SSD) volumes offer storage with consistent and low-latency performance, and are designed for I/O intensive applications such as large relational or NoSQL databases. Magnetic volumes provide the lowest cost per gigabyte of all EBS volume types. Magnetic volumes are ideal for workloads where data is accessed infrequently, and applications where the lowest storage cost is important.

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

NEW QUESTION 77

Do you need to shutdown your EC2 instance when you create a snapshot of EBS volumes that serve as root devices?

- A. No, you only need to shutdown an instance before deleting it.
- B. Yes
- C. No, the snapshot would turn off your instance automatically.
- D. No

Answer: B

Explanation:

Yes, to create a snapshot for Amazon EBS volumes that serve as root devices, you should stop the instance before taking the snapshot.

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

NEW QUESTION 81

Name the disk storage supported by Amazon Elastic Compute Cloud (EC2).

- A. None of these
- B. Amazon AppStream store
- C. Amazon SNS store
- D. Amazon Instance Store

Answer: D

Explanation:

Amazon EC2 supports the following storage options: Amazon Elastic Block Store (Amazon EBS) Amazon EC2 Instance Store Amazon Simple Storage Service (Amazon S3)

Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/Storage.html>

NEW QUESTION 85

A scope has been handed to you to set up a super fast gaming server and you decide that you will use Amazon DynamoDB as your database. For efficient access to data in a table, Amazon DynamoDB creates and maintains indexes for the primary key attributes. A secondary index is a data structure that contains a subset of attributes from a table, along with an alternate key to support Query operations. How many types of secondary indexes does DynamoDB support?

- A. 2
- B. 16
- C. 4
- D. As many as you need

Answer: A

Explanation:

DynamoDB supports two types of secondary indexes:

Local secondary index — an index that has the same hash key as the table, but a different range key. A local secondary index is "local" in the sense that every partition of a local secondary index is scoped to a table partition that has the same hash key.

Global secondary index — an index with a hash and range key that can be different from those on the table. A global secondary index is considered "global" because queries on the index can span all of the data in a table, across all partitions.

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

NEW QUESTION 89

Select the correct statement: Within Amazon EC2, when using Linux instances, the device name

/dev/sda1 is .

- A. reserved for EBS volumes
- B. recommended for EBS volumes
- C. recommended for instance store volumes
- D. reserved for the root device

Answer: D

Explanation:

Within Amazon EC2, when using a Linux instance, the device name /dev/sda1 is reserved for the root device.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/device_naming.html

NEW QUESTION 92

You need to set up security for your VPC and you know that Amazon VPC provides two features that you can use to increase security for your VPC: security groups and network access control lists (ACLs). You have already looked into security groups and you are now trying to understand ACLs. Which statement below is incorrect in relation to ACLs?

- A. Supports allow rules and deny rules.
- B. Is stateful: Return traffic is automatically allowed, regardless of any rules.
- C. Processes rules in number order when deciding whether to allow traffic.
- D. Operates at the subnet level (second layer of defense).

Answer: B

Explanation:

Amazon VPC provides two features that you can use to increase security for your VPC:

Security groups—Act as a firewall for associated Amazon EC2 instances, controlling both inbound and outbound traffic at the instance level

Network access control lists (ACLs)—Act as a firewall for associated subnets, controlling both inbound and outbound traffic at the subnet level

Security groups are stateful: (Return traffic is automatically allowed, regardless of any rules) Network ACLs are stateless: (Return traffic must be explicitly allowed by rules)

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

NEW QUESTION 93

You need to create an Amazon Machine Image (AMI) for a customer for an application which does not appear to be part of the standard AWS AMI template that you can see in the AWS console. What are the alternative possibilities for creating an AMI on AWS?

- A. You can purchase an AMIs from a third party but cannot create your own AMI.
- B. You can purchase an AMIs from a third party or can create your own AMI.
- C. Only AWS can create AMIs and you need to wait till it becomes available.
- D. Only AWS can create AMIs and you need to request them to create one for you

Answer: B

Explanation:

You can purchase an AMIs from a third party, including AMIs that come with service contracts from organizations such as Red Hat. You can also create an AMI and sell it to other Amazon EC2 users. After you create an AMI, you can keep it private so that only you can use it, or you can share it with a specified list of AWS accounts. You can also make your custom AMI public so that the community can use it. Building a safe, secure, usable AMI for public consumption is a fairly straightforward process, if you follow a few simple guidelines.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AMIs.htm>

NEW QUESTION 98

Which one of the below is not an AWS Storage Service?

- A. Amazon S3
- B. Amazon Glacier
- C. Amazon CloudFront
- D. Amazon EBS

Answer: C

Explanation:

AWS Storage Services are: Amazon S3

Amazon Glacier Amazon EBS

AWS Storage Gateway

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

NEW QUESTION 101

You have a lot of data stored in the AWS Storage Gateway and your manager has come to you asking about how the billing is calculated, specifically the Virtual Tape Shelf usage. What would be a correct response to this?

- A. You are billed for the virtual tape data you store in Amazon Glacier and are billed for the size of the virtual tape.
- B. You are billed for the virtual tape data you store in Amazon Glacier and billed for the portion of virtual tape capacity that you use, not for the size of the virtual tape.
- C. You are billed for the virtual tape data you store in Amazon S3 and billed for the portion of virtual tape capacity that you use, not for the size of the virtual tape.
- D. You are billed for the virtual tape data you store in Amazon S3 and are billed for the size of the virtual tape.

Answer: B

Explanation:

The AWS Storage Gateway is a service connecting an on-premises software appliance with cloud-based storage to provide seamless and secure integration between an organization's on-premises IT environment and AWS's storage infrastructure.

AWS Storage Gateway billing is as follows. Volume storage usage (per GB per month):

You are billed for the Cached volume data you store in Amazon S3. You are only billed for volume capacity you use, not for the size of the volume you create.

Snapshot Storage usage (per GB per month): You are billed for the snapshots your gateway stores in Amazon S3. These snapshots are stored and billed as Amazon EBS snapshots. Snapshots are incremental backups, reducing your storage charges. When taking a new snapshot, only the data that has changed since your last snapshot is stored.

Virtual Tape Library usage (per GB per month):

You are billed for the virtual tape data you store in Amazon S3. You are only billed for the portion of virtual tape capacity that you use, not for the size of the virtual tape.

Virtual Tape Shelf usage (per GB per month):

You are billed for the virtual tape data you store in Amazon Glacier. You are only billed for the portion of virtual tape capacity that you use, not for the size of the virtual tape.

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

NEW QUESTION 103

You are configuring a new VPC for one of your clients for a cloud migration project, and only a public VPN will be in place. After you created your VPC, you created a new subnet, a new internet gateway, and attached your internet gateway to your VPC. When you launched your first instance into your VPC, you realized that you aren't able to connect to the instance, even if it is configured with an elastic IP. What should be done to access the instance?

- A. A route should be created as 0.0.0.0/0 and your internet gateway as target.
- B. Attach another ENI to the instance and connect via new ENI.
- C. A NAT instance should be created and all traffic should be forwarded to NAT instance.
- D. A NACL should be created that allows all outbound traffic

Answer: A

Explanation:

All traffic should be routed via Internet Gateway. So, a route should be created with 0.0.0.0/0 as a source, and your Internet Gateway as your target.

Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario1.html

NEW QUESTION 107

You need to create a JSON-formatted text file for AWS CloudFormation. This is your first template and the only thing you know is that the templates include several major sections but there is only one that is required for it to work. What is the only section required?

- A. Mappings
- B. Outputs
- C. Resources
- D. Conditions

Answer: C

Explanation:

AWS CloudFormation is a service that helps you model and set up your Amazon Web Services resources so that you can spend less time managing those resources and more time focusing on your applications that run in AWS. You create a template that describes all the AWS resources that you want (like Amazon EC2 instances or Amazon RDS DB instances), and AWS CloudFormation takes care of provisioning and configuring those resources for you.

A template is a JSON-formatted text file that describes your AWS infrastructure. Templates include several major sections.

The Resources section is the only section that is required.

The first character in the template must be an open brace ({), and the last character must be a closed brace (}). The following template fragment shows the template structure and sections.

Reference: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/template-anatomy.html>

NEW QUESTION 108

You are planning and configuring some EBS volumes for an application. In order to get the most performance out of your EBS volumes, you should attach them to an instance with enough to support your volumes.

- A. Redundancy
- B. Storage
- C. Bandwidth
- D. Memory

Answer: C

Explanation:

When you plan and configure EBS volumes for your application, it is important to consider the configuration of the instances that you will attach the volumes to. In order to get the most performance out of your EBS volumes, you should attach them to an instance with enough bandwidth to support your volumes, such as an EBS-optimized instance or an instance with 10 Gigabit network connectivity. This is especially important when you use General Purpose (SSD) or Provisioned IOPS (SSD) volumes, or when you stripe multiple volumes together in a RAID configuration.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-ec2-config.html>

NEW QUESTION 112

Can a single EBS volume be attached to multiple EC2 instances at the same time?

- A. Yes
- B. No
- C. Only for high-performance EBS volumes.
- D. Only when the instances are located in the US region

Answer: B

Explanation:

You can't attach an EBS volume to multiple EC2 instances. This is because it is equivalent to using a single hard drive with many computers at the same time.
Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonEBS.html>

NEW QUESTION 116

You have been setting up an Amazon Virtual Private Cloud (Amazon VPC) for your company, including setting up subnets. Security is a concern, and you are not sure which is the best security practice for securing subnets in your VPC. Which statement below is correct in describing the protection of AWS resources in each subnet?

- A. You can use multiple layers of security, including security groups and network access control lists (ACL).
- B. You can only use access control lists (ACL).
- C. You don't need any security in subnets.
- D. You can use multiple layers of security, including security groups, network access control lists (ACL) and CloudHSM.

Answer: A

Explanation:

A subnet is a range of IP addresses in your VPC. You can launch AWS resources into a subnet that you select. Use a public subnet for resources that must be connected to the Internet, and a private subnet for resources that won't be connected to the Internet.

To protect the AWS resources in each subnet, you can use multiple layers of security, including security groups and network access control lists (ACL).

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

NEW QUESTION 119

Your organization is in the business of architecting complex transactional databases. For a variety of reasons, this has been done on EBS. What is AWS's recommendation for customers who have architected databases using EBS for backups?

- A. Backups to Amazon S3 be performed through the database management system.
- B. Backups to AWS Storage Gateway be performed through the database management system.
- C. If you take regular snapshots no further backups are required.
- D. Backups to Amazon Glacier be performed through the database management system.

Answer: A

Explanation:

Data stored in Amazon EBS volumes is redundantly stored in multiple physical locations as part of normal operation of those services and at no additional charge. However, Amazon EBS replication is stored within the same availability zone, not across multiple zones; therefore, it is highly recommended that you conduct regular snapshots to Amazon S3 for long-term data durability.

For customers who have architected complex transactional databases using EBS, it is recommended that backups to Amazon S3 be performed through the database management system so that distributed transactions and logs can be checkpointed.

AWS does not perform backups of data that are maintained on virtual disks attached to running instances on Amazon EC2.

Reference: <http://d0.awsstatic.com/whitepapers/Security/AWS%20Security%20Whitepaper.pdf>

NEW QUESTION 124

A user has created a CloudFormation stack. The stack creates AWS services, such as EC2 instances, ELB, AutoScaling, and RDS. While creating the stack it created EC2, ELB and AutoScaling but failed to create RDS. What will CloudFormation do in this scenario?

- A. Rollback all the changes and terminate all the created services
- B. It will wait for the user's input about the error and correct the mistake after the input
- C. CloudFormation can never throw an error after launching a few services since it verifies all the steps before launching
- D. It will warn the user about the error and ask the user to manually create RDS

Answer: A

Explanation:

AWS CloudFormation is an application management tool which provides application modeling, deployment, configuration, management and related activities. The AWS CloudFormation stack is a collection of AWS resources which are created and managed as a single unit when AWS CloudFormation instantiates a template. If any of the services fails to launch, CloudFormation will rollback all the changes and terminate or delete all the created services.

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

NEW QUESTION 127

You're trying to delete an SSL certificate from the IAM certificate store, and you're getting the message "Certificate: <certificate-id> is being used by CloudFront." Which of the following statements is probably the reason why you are getting this error?

- A. Before you can delete an SSL certificate, you need to either rotate SSL certificates or revert from using a custom SSL certificate to using the default CloudFront certificate.
- B. You can't delete SSL certificates. You need to request it from AWS.
- C. Before you can delete an SSL certificate, you need to set up the appropriate access level in IAM
- D. Before you can delete an SSL certificate you need to set up https on your server

Answer: A

Explanation:

CloudFront is a web service that speeds up distribution of your static and dynamic web content, for example, .html, .css, .php, and image files, to end users.

Every CloudFront web distribution must be associated either with the default CloudFront certificate or with a custom SSL certificate. Before you can delete an SSL certificate, you need to either rotate SSL certificates (replace the current custom SSL certificate with another custom SSL certificate) or revert from using a custom SSL certificate to using the default CloudFront certificate.

Reference: <http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Troubleshooting.htm>

NEW QUESTION 130

How many types of block devices does Amazon EC2 support?

- A. 4
- B. 5
- C. 2
- D. 1

Answer: C

Explanation:

Amazon EC2 supports 2 types of block devices. Reference:

<http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/block-device-mapping-concepts.html>

NEW QUESTION 133

Any person or application that interacts with AWS requires security credentials. AWS uses these credentials to identify who is making the call and whether to allow the requested access. You have just set up a VPC network for a client and you are now thinking about the best way to secure this network. You set up a security group called vpcsecuritygroup. Which following statement is true in respect to the initial settings that will be applied to this security group if you choose to use the default settings for this group?

- A. Allow all inbound traffic and allow no outbound traffic.
- B. Allow no inbound traffic and allow all outbound traffic.
- C. Allow inbound traffic on port 80 only and allow all outbound traffic.
- D. Allow all inbound traffic and allow all outbound traffic

Answer: B

Explanation:

Amazon VPC provides advanced security features such as security groups and network access control lists to enable inbound and outbound filtering at the instance level and subnet level.

AWS assigns each security group a unique ID in the form sg-xxxxxxx. The following are the initial settings for a security group that you create:

Allow no inbound traffic Allow all outbound traffic

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html>

NEW QUESTION 134

You are using Amazon SES as an email solution but are unsure of what its limitations are. Which statement below is correct in regards to that?

- A. New Amazon SES users who have received production access can send up to 1,000 emails per 24-hour period, at a maximum rate of 10 emails per second.
- B. Every Amazon SES sender has the same set of sending limits
- C. Sending limits are based on messages rather than on recipients
- D. Every Amazon SES sender has a unique set of sending limits

Answer: D

Explanation:

Amazon Simple Email Service (Amazon SES) is a highly scalable and cost-effective email-sending service for businesses and developers. Amazon SES eliminates the complexity and expense of building an in-house email solution or licensing, installing, and operating a third-party email service for this type of email communication.

Every Amazon SES sender has a unique set of sending limits, which are calculated by Amazon SES on an ongoing basis:

Sending quota — the maximum number of emails you can send in a 24-hour period. Maximum send rate — the maximum number of emails you can send per second.

New Amazon SES users who have received production access can send up to 10,000 emails per 24-hour period, at a maximum rate of 5 emails per second.

Amazon SES automatically adjusts these limits upward, as long as you send high-quality email. If your existing quota is not adequate for your needs and the system has not automatically increased your quota, you can submit an SES Sending Quota Increase case at any time.

Sending limits are based on recipients rather than on messages. You can check your sending limits at any time by using the Amazon SES console.

Note that if your email is detected to be of poor or QUESTION able quality (e.g., high complaint rates, high bounce rates, spam, or abusive content), Amazon SES might temporarily or permanently reduce your permitted send volume, or take other action as AWS deems appropriate.

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

NEW QUESTION 139

You have been asked to build AWS infrastructure for disaster recovery for your local applications and within that you should use an AWS Storage Gateway as part of the solution. Which of the following best describes the function of an AWS Storage Gateway?

- A. Accelerates transferring large amounts of data between the AWS cloud and portable storage devices .
- B. A web service that speeds up distribution of your static and dynamic web content.
- C. Connects an on-premises software appliance with cloud-based storage to provide seamless and secure integration between your on-premises IT environment and AWS's storage infrastructure.
- D. Is a storage service optimized for infrequently used data, or "cold data."

Answer: C

Explanation:

AWS Storage Gateway connects an on-premises software appliance with cloud-based storage to provide seamless integration with data security features between your on-premises IT environment and the Amazon Web Services (AWS) storage infrastructure. You can use the service to store data in the AWS cloud for scalable and cost-effective storage that helps maintain data security. AWS Storage Gateway offers both volume-based and tape-based storage solutions:

Volume gateways Gateway-cached volumes Gateway-stored volumes

Gateway-virtual tape library (VTL)

Reference:

http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_disasterrecovery_07.pdf

NEW QUESTION 144

In Amazon CloudFront, if you use Amazon EC2 instances and other custom origins with CloudFront, it is recommended to .

- A. not use Elastic Load Balancing
- B. restrict Internet communication to private instances while allowing outgoing traffic
- C. enable access key rotation for CloudWatch metrics
- D. specify the URL of the load balancer for the domain name of your origin server

Answer: D

Explanation:

In Amazon CloudFront, you should use an Elastic Load Balancing load balancer to handle traffic across multiple Amazon EC2 instances and to isolate your application from changes to Amazon EC2 instances. When you create your CloudFront distribution, specify the URL of the load balancer for the domain name of your origin server.

Reference: <http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/CustomOriginBestPractices.html>

NEW QUESTION 146

A user is sending bulk emails using AWS SES. The emails are not reaching some of the targeted audience because they are not authorized by the ISPs. How can the user ensure that the emails are all delivered?

- A. Send an email using DKIM with SES.
- B. Send an email using SMTP with SES.
- C. Open a ticket with AWS support to get it authorized with the ISP.
- D. Authorize the ISP by sending emails from the development account

Answer: A

Explanation:

Domain Keys Identified Mail (DKIM) is a standard that allows senders to sign their email messages and ISPs, and use those signatures to verify that those messages are legitimate and have not been modified by a third party in transit.

Reference: <http://docs.aws.amazon.com/ses/latest/DeveloperGuide/dkim.html>

NEW QUESTION 150

You are setting up your first Amazon Virtual Private Cloud (Amazon VPC) network so you decide you should probably use the AWS Management Console and the VPC Wizard. Which of the following is not an option for network architectures after launching the "Start VPC Wizard" in Amazon VPC page on the AWS Management Console?

- A. VPC with a Single Public Subnet Only
- B. VPC with a Public Subnet Only and Hardware VPN Access
- C. VPC with Public and Private Subnets and Hardware VPN Access
- D. VPC with a Private Subnet Only and Hardware VPN Access

Answer: B

Explanation:

Amazon VPC enables you to build a virtual network in the AWS cloud - no VPNs, hardware, or physical datacenters required.

Your AWS resources are automatically provisioned in a ready-to-use default VPC. You can choose to create additional VPCs by going to Amazon VPC page on the AWS Management Console and click on the "Start VPC Wizard" button.

You'll be presented with four basic options for network architectures. After selecting an option, you can modify the size and IP address range of the VPC and its subnets. If you select an option with Hardware VPN Access, you will need to specify the IP address of the VPN hardware on your network. You can modify the VPC to add more subnets or add or remove gateways at any time after the VPC has been created.

The four options are:

VPC with a Single Public Subnet Only VPC with Public and Private Subnets

VPC with Public and Private Subnets and Hardware VPN Access VPC with a Private Subnet Only and Hardware VPN Access Reference:

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

NEW QUESTION 152

An EC2 instance is connected to an ENI (Elastic Network Interface) in one subnet. What happens when you attach an ENI of a different subnet to this EC2 instance?

- A. The EC2 instance follows the rules of the older subnet
- B. The EC2 instance follows the rules of both the subnets
- C. Not possible, cannot be connected to 2 ENIs
- D. The EC2 instance follows the rules of the newer subnet

Answer: B

Explanation:

AWS allows you create an elastic network interface (ENI), attach an ENI to an EC2 instance, detach an ENI from an EC2 instance and attach this ENI to another EC2 instance. The attributes of a network traffic follow the ENI which is attached to an EC2 instance or detached from an EC2 instance. When you move an ENI from one EC2 instance to another, network traffic is redirected to the new EC2 instance. You can create and attach additional ENIs to an EC2 instance.

Attaching multiple network interfaces (ENIs) to an EC2 instance is useful to: Create a management network.

Use network and security appliances in your VPC.

Create dual-homed instances with workloads/roles on distinct subnets Create a low-budget, high-availability solution.

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

NEW QUESTION 154

A user is trying to launch a similar EC2 instance from an existing instance with the option "Launch More like this". The AMI of the selected instance is deleted. What will happen in this case?

- A. AWS does not need an AMI for the "Launch more like this" option
- B. AWS will launch the instance but will not create a new AMI
- C. AWS will create a new AMI and launch the instance
- D. AWS will throw an error saying that the AMI is deregistered

Answer: D

Explanation:

If the user has deregistered the AMI of an EC2 instance and is trying to launch a similar instance with the option "Launch more like this", AWS will throw an error saying that the AMI is deregistered or not available.

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

NEW QUESTION 159

A company wants to review the security requirements of Glacier. Which of the below mentioned statements is true with respect to the AWS Glacier data security?

- A. All data stored on Glacier is protected with AES-256 serverside encryption.
- B. All data stored on Glacier is protected with AES-128 serverside encryption.
- C. The user can set the serverside encryption flag to encrypt the data stored on Glacier.
- D. The data stored on Glacier is not encrypted by default

Answer: A

Explanation:

For Amazon Web Services, all the data stored on Amazon Glacier is protected using serverside encryption. AWS generates separate unique encryption keys for each Amazon Glacier archive, and encrypts it using AES-256. The encryption key then encrypts itself using AES-256 with a master key that is stored in a secure location.

Reference: http://media.amazonwebservices.com/AWS_Security_Best_Practices.pdf

NEW QUESTION 163

You are architecting a highly-scalable and reliable web application which will have a huge amount of content. You have decided to use CloudFront as you know it will speed up distribution of your static and dynamic web content and know that Amazon CloudFront integrates with Amazon CloudWatch metrics so that you can monitor your web application. Because you live in Sydney you have chosen the the Asia Pacific (Sydney) region in the AWS console. However you have set up this up but no CloudFront metrics seem to be appearing in the CloudWatch console. What is the most likely reason from the possible choices below for this?

- A. Metrics for CloudWatch are available only when you choose the same region as the application you are monitoring.
- B. You need to pay for CloudWatch for it to become active.
- C. Metrics for CloudWatch are available only when you choose the US East (Virginia)
- D. Metrics for CloudWatch are not available for the Asia Pacific region as yet

Answer: C

Explanation:

CloudFront is a global service, and metrics are available only when you choose the US East (N. Virginia) region in the AWS console. If you choose another region, no CloudFront metrics will appear in the CloudWatch console.

Reference:

<http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/monitoring-using-cloudwatch.html>

NEW QUESTION 168

After a major security breach your manager has requested a report of all users and their credentials in AWS. You discover that in IAM you can generate and download a credential report that lists all users in your account and the status of their various credentials, including passwords, access keys, MFA devices, and signing certificates. Which following statement is incorrect in regards to the use of credential reports?

- A. Credential reports are downloaded XML files.
- B. You can get a credential report using the AWS Management Console, the AWS CLI, or the IAM API.
- C. You can use the report to audit the effects of credential lifecycle requirements, such as password rotation.
- D. You can generate a credential report as often as once every four hour

Answer: A

Explanation:

To access your AWS account resources, users must have credentials.

You can generate and download a credential report that lists all users in your account and the status of their various credentials, including passwords, access keys, MFA devices, and signing certificates. You can get a credential report using the AWS Management Console, the AWS CLI, or the IAM API.

You can use credential reports to assist in your auditing and compliance efforts. You can use the report to audit the effects of credential lifecycle requirements, such as password rotation. You can provide the report to an external auditor, or grant permissions to an auditor so that he or she can download the report directly.

You can generate a credential report as often as once every four hours. When you request a report, IAM first checks whether a report for the account has been generated within the past four hours. If so, the most recent report is downloaded. If the most recent report for the account is more than four hours old, or if there are no previous reports for the account, IAM generates and downloads a new report.

Credential reports are downloaded as comma-separated values (CSV) files.

You can open CSV files with common spreadsheet software to perform analysis, or you can build an application that consumes the CSV files programmatically and performs custom analysis. Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/credential-reports.html>

NEW QUESTION 171

A user is planning a highly available application deployment with EC2. Which of the below mentioned options will not help to achieve HA?

- A. Elastic IP address
- B. PIOPS
- C. AMI
- D. Availability Zones

Answer: B

Explanation:

In Amazon Web Service, the user can achieve HA by deploying instances in multiple zones. The elastic IP helps the user achieve HA when one of the instances is down but still keeps the same URL. The AMI helps launching the new instance. The PIOPS is for the performance of EBS and does not help for HA. Reference: http://media.amazonwebservices.com/AWS_Web_Hosting_Best_Practices.pdf

NEW QUESTION 172

Which DNS name can only be resolved within Amazon EC2?

- A. Public DNS name
- B. Internal DNS name
- C. External DNS name
- D. Global DNS name

Answer: B

Explanation:

Only Internal DNS name can be resolved within Amazon EC2. Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/using-instance-addressing.html>

NEW QUESTION 176

You need to create a management network using network interfaces for a virtual private cloud (VPC) network. Which of the following statements is incorrect pertaining to Best Practices for Configuring Network Interfaces.

- A. You can detach secondary (ethN) network interfaces when the instance is running or stoppe
- B. However, you can't detach the primary (eth0) interface.
- C. Launching an instance with multiple network interfaces automatically configures interfaces, private IP addresses, and route tables on the operating system of the instance.
- D. You can attach a network interface in one subnet to an instance in another subnet in the same VPC, however, both the network interface and the instance must reside in the same Availability Zone.
- E. Attaching another network interface to an instance is a valid method to increase or double the network bandwidth to or from the dual-homed instance

Answer: D

Explanation:

Best Practices for Configuring Network Interfaces

You can attach a network interface to an instance when it's running (hot attach), when it's stopped (warm attach), or when the instance is being launched (cold attach).

You can detach secondary (ethN) network interfaces when the instance is running or stopped. However, you can't detach the primary (eth0) interface.

You can attach a network interface in one subnet to an instance in another subnet in the same VPC, however, both the network interface and the instance must reside in the same Availability Zone.

When launching an instance from the CLI or API, you can specify the network interfaces to attach to the instance for both the primary (eth0) and additional network interfaces.

Launching an instance with multiple network interfaces automatically configures interfaces, private IP addresses, and route tables on the operating system of the instance.

A warm or hot attach of an additional network interface may require you to manually bring up the second interface, configure the private IP address, and modify the route table accordingly. (Instances running Amazon Linux automatically recognize the warm or hot attach and configure themselves.)

Attaching another network interface to an instance is not a method to increase or double the network bandwidth to or from the dual-homed instance.

Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html#use-network-and-security-apliances-in-your-vpc>

NEW QUESTION 180

You can seamlessly join an EC2 instance to your directory domain. What connectMty do you need to be able to connect remotely to this instance?

- A. You must have IP connectMty to the instance from the network you are connecting from.
- B. You must have the correct encryption keys to connect to the instance remotely.
- C. You must have enough bandwidth to connect to the instance.
- D. You must use MFA authentication to be able to connect to the instance remotel

Answer: A

Explanation:

You can seamlessly join an EC2 instance to your directory domain when the instance is launched using the Amazon EC2 Simple Systems Manager. If you need to manuallyjoin an EC2 instance to your domain, you must launch the instance in the proper region and security group or subnet, then join the instance to the domain. To be able to connect remotely to these instances, you must have IP connectMty to the instances from the network you are connecting from. In most cases, this requires that an Internet gateway be attached to your VPC and that the instance has a public IP address.

Reference: http://docs.aws.amazon.com/directoryservice/latest/admin-guide/join_a_directory.html

NEW QUESTION 181

You are in the process of moving your friend's WordPress site onto AWS to try and save him some money, and you have told him that he should probably also move his domain name. He asks why he can't leave his domain name where it is and just have his infrastructure on AWS. What would be an incorrect response to his question ?

- A. Route 53 offers low query latency for your end users.
- B. Route 53 is designed to automatically answer queries from the optimal location depending on network conditions.
- C. The globally distributed nature of AWS's DNS servers helps ensure a consistent ability to route your end users to your application.
- D. Route 53 supports Domain Name System Security Extensions (DNSSEC).

Answer: D

Explanation:

Amazon Route 53 provides highly available and scalable Domain Name System (DNS), domain name registration, and health-checking web services. Route 53 is built using AWS's highly available and reliable infrastructure. The globally distributed nature of our DNS servers helps ensure a consistent ability to route your end users to your application by circumventing any internet or network related issues. Route 53 is designed to provide the level of dependability required by important applications. Using a global anycast network of DNS servers around the world, Route 53 is designed to automatically answer queries from the optimal location depending on network conditions. As a result, the service offers low query latency for your end users. Amazon Route 53 does not support Domain Name System Security Extensions (DNSSEC) at this time. Reference: <https://aws.amazon.com/route53/faqs/>

NEW QUESTION 184

In Amazon EC2, you are billed instance-hours when .

- A. your EC2 instance is in a running state
- B. the instance exits from Amazon S3 console
- C. your instance still exits the EC2 console
- D. EC2 instances stop

Answer: A

Explanation:

You are billed instance-hours as long as your EC2 instance is in a running state. Reference: <http://aws.amazon.com/ec2/faqs/>

NEW QUESTION 188

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 Email Service.
- C. AWS NMobile Communication Service.
- D. AWS Simple Queue 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 189

You have written a CloudFormation template that creates 1 Elastic Load Balancer fronting 2 EC2 Instances. Which section of the template should you edit so that the DNS of the load balancer is returned upon creation of the stack?

- A. Resources
- B. Outputs
- C. Parameters
- D. Mappings

Answer: B

Explanation:

You can use AWS CloudFormation's sample templates or create your own templates to describe the AWS resources, and any associated dependencies or runtime parameters, required to run your application.

Reference:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/outputs-section-structure.html>

NEW QUESTION 194

How can you apply more than 100 rules to an Amazon EC2-Classic?

- A. By adding more security groups
- B. You need to create a default security group specifying your required rules if you need to use more than 100 rules per security group.
- C. By default the Amazon EC2 security groups support 500 rules.
- D. You can't add more than 100 rules to security groups for an Amazon EC2 instanc

Answer: D

Explanation:

In EC2-Classic, you can associate an instance with up to 500 security groups and add up to 100 rules to a security group.

Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/using-network-security.html>

NEW QUESTION 198

You have a Business support plan with AWS. One of your EC2 instances is running Microsoft Windows Server 2008 R2 and you are having problems with the

software. Can you receive support from AWS for this software?

- A. Yes
- B. No, AWS does not support any third-party software.
- C. No, Microsoft Windows Server 2008 R2 is not supported.
- D. No, you need to be on the enterprise support plan

Answer: A

Explanation:

Third-party software support is available only to AWS Support customers enrolled for Business or Enterprise Support. Third-party support applies only to software running on Amazon EC2 and does not extend to assisting with on-premises software. An exception to this is a VPN tunnel configuration running supported devices for Amazon VPC.

Reference: <https://aws.amazon.com/premiumsupport/features/>

NEW QUESTION 200

In Amazon EC2, how many Elastic IP addresses can you have by default?

- A. 10
- B. 2
- C. 5
- D. 20

Answer: C

Explanation:

The number of Elastic IP addresses you can have in EC2 is 5.

Reference: http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#limits_ec2

NEW QUESTION 204

After deciding that EMR will be useful in analysing vast amounts of data for a gaming website that you are architecting you have just deployed an Amazon EMR Cluster and wish to monitor the cluster performance. Which of the following tools cannot be used to monitor the cluster performance?

- A. Kinesis
- B. Ganglia
- C. CloudWatch Metrics
- D. Hadoop Web Interfaces

Answer: A

Explanation:

Amazon EMR provides several tools to monitor the performance of your cluster. Hadoop Web Interfaces

Every cluster publishes a set of web interfaces on the master node that contain information about the cluster. You can access these web pages by using an SSH tunnel to connect them on the master node. For more information, see [View Web Interfaces Hosted on Amazon EMR Clusters](#).

CloudWatch Metrics

Every cluster reports metrics to CloudWatch. CloudWatch is a web service that tracks metrics, and which you can use to set alarms on those metrics. For more information, see [Monitor Metrics with CloudWatch](#). Ganglia

Ganglia is a cluster monitoring tool. To have this available, you have to install Ganglia on the cluster when you launch it. After you've done so, you can monitor the cluster as it runs by using an SSH tunnel to connect to the Ganglia UI running on the master node. For more information, see [Monitor Performance with Ganglia](#).

Reference:

<http://docs.aws.amazon.com/ElasticMapReduce/latest/DeveloperGuide/emr-troubleshoot-tools.html>

NEW QUESTION 206

Can you move a Reserved Instance from one Availability Zone to another?

- A. Yes, but each Reserved Instance is associated with a specific Region that cannot be changed.
- B. Yes, only in US-West-2.
- C. Yes, only in US-East-1.
- D. No

Answer: A

Explanation:

Each Reserved Instance is associated with a specific Region, which is fixed for the lifetime of the reservation and cannot be changed. Each reservation can, however, be used in any of the available AZs within the associated Region.

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

NEW QUESTION 208

An application hosted at the EC2 instance receives an HTTP request from ELB. The same request has an X-Forwarded-For header, which has three IP addresses. Which system's IP will be a part of this header?

- A. Previous Request IP address.
- B. Client IP address.
- C. All of the answers listed here.
- D. Load Balancer IP address

Answer: C

Explanation:

When a user sends a request to ELB over HTTP/HTTPS, the request header log at the instance will only receive the IP of ELB. This is because ELB is the interceptor between the EC2 instance and the client request. To get the client IP, use the header X-Forwarded-For in header. The client IP address in the X-Forwarded-For request header is followed by the IP addresses of each successive proxy that passes along the request. The last IP address is the IP address that connects to the back-end application instance. e.g. if the HTTP request already has a header when it reaches the Load Balancer, the IP address from which the request came is appended at the end of the header followed by the IP address of the Load Balancer. In such cases, the X-Forwarded-For request header takes the following form:

X-Forwarded-For: clientIPAddress, previousRequestIPAddress, LoadBalancerIPAddress. Reference:

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

NEW QUESTION 213

You need to set up a security certificate for a client's e-commerce website as it will use the HTTPS protocol. Which of the below AWS services do you need to access to manage your SSL server certificate?

- A. AWS Directory Service
- B. AWS Identity & Access Management
- C. AWS CloudFormation
- D. Amazon Route 53

Answer: B

Explanation:

AWS Identity and Access Management (IAM) is a web service that enables Amazon Web Services (AWS) customers to manage users and user permissions in AWS.

All your SSL server certificates are managed by AWS Identity and Access management (IAM). Reference:

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

NEW QUESTION 217

Company B is launching a new game app for mobile devices. Users will log into the game using their existing social media account to streamline data capture. Company B would like to directly save player data and scoring information from the mobile app to a DynamoDB table named Score Data

When a user saves their game the progress data will be stored to the Game state 53 bucket. What is the best approach for storing data to DynamoDB and 53?

- A. Use an EC2 Instance that is launched with an EC2 role providing access to the Score Data DynamoDB table and the GameState 53 bucket that communicates with the mobile app via web services.
- B. Use temporary security credentials that assume a role providing access to the Score Data DynamoDB table and the Game State 53 bucket using web identity federation.
- C. Use Login with Amazon allowing users to sign in with an Amazon account providing the mobile app with access to the Score Data DynamoDB table and the Game State 53 bucket.
- D. Use an IAM user with access credentials assigned a role providing access to the Score Data DynamoDB table and the Game State 53 bucket for distribution with the mobile app.

Answer: B

Explanation:

Web Identity Federation

Imagine that you are creating a mobile app that accesses AWS resources, such as a game that runs on a mobile device and stores player and score information using Amazon S3 and DynamoDB. When you write such an app, you'll make requests to AWS services that must be signed with an AWS access key. However, we strongly recommend that you do not embed or distribute long-term AWS credentials with apps that a user downloads to a device, even in an encrypted store. Instead, build your app so that it requests temporary AWS security credentials dynamically when needed using web identity federation. The supplied temporary credentials map to an AWS role that has only the permissions needed to perform the tasks required by the mobile app.

With web identity federation, you don't need to create custom sign-in code or manage your own user identities. Instead, users of your app can sign in using a well-known identity provider (IdP) - such as Login with Amazon, Facebook, Google, or any other OpenID Connect (OIDC)-compatible IdP, receive an authentication token, and then exchange that token for temporary security credentials in AWS that map to an IAM role with permissions to use the resources in your AWS account. Using an IdP helps you keep your AWS account secure, because you don't have to embed and distribute long-term security credentials with your application.

For most scenarios, we recommend that you use Amazon Cognito because it acts as an identity broker and does much of the federation work for you. For details, see the following section, [Using Amazon Cognito for Mobile Apps](#).

If you don't use Amazon Cognito, then you must write code that interacts with a web IdP (Login with Amazon, Facebook, Google, or any other OIDC-compatible IdP) and then calls the Assume Role With Web Identity API to trade the authentication token you get from those IdPs for AWS temporary security credentials. If you have already used this approach for existing apps, you can continue to use it.

Using Amazon Cognito for Mobile Apps

The preferred way to use web identity federation is to use Amazon Cognito. For example, Adele the developer is building a game for a mobile device where user data such as scores and profiles is stored in Amazon S3 and Amazon DynamoDB. Adele could also store this data locally on the device and use Amazon Cognito to keep it synchronized across devices. She knows that for security and maintenance reasons, long-term AWS security credentials should not be distributed with the game. She also knows that the game might have a large number of users. For all of these reasons, she does not want to create new user identities in IAM for each player. Instead, she builds the game so that users can sign in using an identity that they've already established with a well-known identity provider, such as Login with Amazon, Facebook, Google, or any OpenID Connect (OIDC)-compatible identity provider.

Her game can take advantage of the authentication mechanism from one of these providers to validate the user's identity.

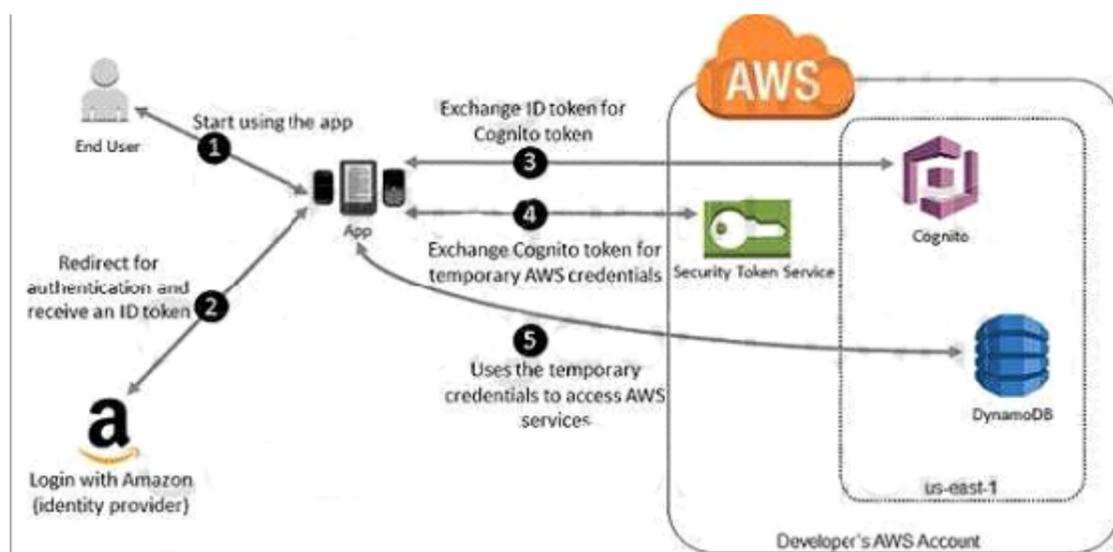
To enable the mobile app to access her AWS resources, Adele first registers for a developer ID with her chosen IdPs. She also configures the application with each of these providers. In her AWS account that contains the Amazon S3 bucket and DynamoDB table for the game, Adele uses Amazon Cognito to create IAM roles that precisely define permissions that the game needs. If she is using an OIDC IdP, she also creates an IAM OIDC identity provider entity to establish trust between her AWS account and the IdP.

In the app's code, Adele calls the sign-in interface for the IdP that she configured previously. The IdP handles all the details of letting the user sign in, and the app gets an OAuth access token or OIDC ID token from the provider. Adele's app can trade this authentication information for a set of temporary security credentials that consist of an AWS access key ID, a secret access key, and a session token.

The app can then use these credentials to access web services offered by AWS. The app is limited to the permissions that are defined in the role that it assumes. The following figure shows a simplified flow for how this might work, using Login with Amazon as the IdP.

For Step 2, the app can also use Facebook, Google, or any OIDC-compatible identity provider, but that's not shown here.

Sample workflow using Amazon Cognito to federate users for a mobile application



A customer starts your app on a mobile device. The app asks the user to sign in. The app uses Login with Amazon resources to accept the user's credentials. The app uses Cognito APIs to exchange the Login with Amazon ID token for a Cognito token. The app requests temporary security credentials from AWS STS, passing the Cognito token.

The temporary security credentials can be used by the app to access any AWS resources required by the app to operate. The role associated with the temporary security credentials and its assigned policies determines what can be accessed.

Use the following process to configure your app to use Amazon Cognito to authenticate users and give your app access to AWS resources. For specific steps to accomplish this scenario, consult the documentation for Amazon Cognito.

(Optional) Sign up as a developer with Login with Amazon, Facebook, Google, or any other OpenID Connect (OIDC)-compatible identity provider and configure one or more apps with the provider. This step is optional because Amazon Cognito also supports unauthenticated (guest) access for your users.

Go to Amazon Cognito in the AWS IAM console. Use the Amazon Cognito wizard to create an identity pool, which is a container that Amazon Cognito uses to keep end user identities organized for your apps. You can share identity pools between apps. When you set up an identity pool, Amazon Cognito creates one or two IAM roles (one for authenticated identities, and one for unauthenticated "guest" identities) that define permissions for Amazon Cognito users.

Download and integrate the AWS SDK for iOS or the AWS SDK for Android with your app, and import the files required to use Amazon Cognito.

Create an instance of the Amazon Cognito credentials provider, passing the identity pool ID, your AWS account number, and the Amazon Resource Name (ARN) of the roles that you associated with the identity pool. The Amazon Cognito wizard in the AWS Management Console provides sample code to help you get started.

When your app accesses an AWS resource, pass the credentials provider instance to the client object, which passes temporary security credentials to the client. The permissions for the credentials are based on the role or roles that you defined earlier.

NEW QUESTION 222

You have launched an EC2 instance with four (4) 500GB EBS Provisioned IOPS volumes attached. The EC2 instance is EBS-Optimized and supports 500 Mbps throughput between EC2 and EBS. The two EBS volumes are configured as a single RAID 0 device, and each Provisioned IOPS volume is provisioned with 4,000 IOPS (4,000 16KB reads or writes) for a total of 16,000 random IOPS on the instance. The EC2 instance initially delivers the expected 16,000 IOPS random read and write performance. Sometime later in order to increase the total random I/O performance of the instance, you add an additional two 500 GB EBS Provisioned IOPS volumes to the RAID. Each volume is provisioned to 4,000 IOPS like the original four for a total of 24,000 IOPS on the EC2 instance. Monitoring shows that the EC2 instance CPU utilization increased from 50% to 70%, but the total random IOPS measured at the instance level does not increase at all. What is the problem and a valid solution?

- A. Larger storage volumes support higher Provisioned IOPS rates: increase the provisioned volumestorage of each of the 6 EBS volumes to 1TB
- B. The EBS-Optimized throughput limits the total IOPS that can be utilized: use an EBS-Optimized instance that provides larger throughput.
- C. Small block sizes cause performance degradation, limiting the I/O throughput, configure the instance device driver and file system to use 64KB blocks to increase throughput.
- D. RAID 0 only scales linearly to about 4 devices, use RAID 0 with 4 EBS Provisioned IOPS volumes but increase each Provisioned IOPS EBS volume to 6,000 IOPS.
- E. The standard EBS instance root volume limits the total IOPS rate, change the instance root volume to also be a 500GB 4,000 Provisioned IOPS volume.

Answer: E

NEW QUESTION 225

A web design company currently runs several FTP servers that their 250 customers use to upload and download large graphic files. They wish to move this system to AWS to make it more scalable, but they wish to maintain customer privacy and keep costs to a minimum. What AWS architecture would you recommend?

- A. Ask their customers to use an S3 client instead of an FTP client.
- B. Create a single S3 bucket. Create an IAM user for each customer. Put the IAM users in a group that has an IAM policy that permits access to sub-directories within the bucket via use of the 'username' Policy variable.
- C. Create a single S3 bucket with Reduced Redundancy Storage turned on and ask their customers to use an S3 client instead of an FTP client. Create a bucket for each customer with a Bucket Policy that permits access only to that one customer.
- D. Create an auto-scaling group of FTP servers with a scaling policy to automatically scale-in when minimum network traffic on the auto-scaling group is below a given threshold.
- E. Load a central list of ftp users from S3 as part of the user data startup script on each instance.
- F. Create a single S3 bucket with Requester Pays turned on and ask their customers to use an S3 client instead of an FTP client. Create a bucket for each customer with a Bucket Policy that permits access only to that one customer.

Answer: A

NEW QUESTION 230

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. Route 53 Record Sets
- B. IAM Roles
- C. Elastic IP Addresses (EIP)

- D. EC2 Key Pairs
- E. Launch configurations
- F. Security Groups

Answer: AC

Explanation:

Reference:

http://tech.com/wp-content/themes/optimize/download/AWSDisaster_Recovery.pdf (page 6)

NEW QUESTION 234

Your company runs a customer facing event registration site. This site is built with a 3-tier architecture with web and application tier servers and a MySQL database. The application requires 6 web tier servers and 6 application tier servers for normal operation, but can run on a minimum of 65% server capacity and a single MySQL database. When deploying this application in a region with three availability zones (AZs) which architecture provides high availability?

- A. A web tier deployed across 2 AZs with 3 EC2 (Elastic Compute Cloud) instances in each AZ inside an Auto Scaling Group behind an ELB (elastic load balancer), and an application tier deployed across 2 AZs with 3 EC2 instances in each AZ inside an Auto Scaling Group behind an ELB, and one RDS (Relational Database Service) instance deployed with read replicas in the other AZ.
- B. A web tier deployed across 3 AZs with 2 EC2 (Elastic Compute Cloud) instances in each AZ inside an Auto Scaling Group behind an ELB (elastic load balancer) and an application tier deployed across 3 AZs with 2 EC2 instances in each AZ inside an Auto Scaling Group behind an ELB and one RDS (Relational Database Service) Instance deployed with read replicas in the two other AZs.
- C. A web tier deployed across 2 AZs with 3 EC2 (Elastic Compute Cloud) instances in each AZ inside an Auto Scaling Group behind an ELB (elastic load balancer) and an application tier deployed across 2 AZs with 3 EC2 instances in each AZ inside an Auto Scaling Group behind an ELB and a Multi-AZ RDS (Relational Database Service) deployment.
- D. A web tier deployed across 3 AZs with 2 EC2 (Elastic Compute Cloud) instances in each AZ inside an Auto Scaling Group behind an ELB (elastic load balancer). And an application tier deployed across 3 AZs with 2 EC2 instances in each AZ inside an Auto Scaling Group behind an ELB and a Multi-AZ RDS (Relational Database services) deployment.
- E. And a Multi-AZ RDS (Relational Database services) deployment.

Answer: D

Explanation:

Amazon RDS Multi-AZ Deployments

Amazon RDS Multi-AZ deployments provide enhanced availability and durability for Database (DB) Instances, making them a natural fit for production database workloads. When you provision a Multi-AZ DB Instance, Amazon RDS automatically creates a primary DB Instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). Each AZ runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. In case of an infrastructure failure (for example, instance hardware failure, storage failure, or network disruption), Amazon RDS performs an automatic failover to the standby, so that you can resume database operations as soon as the failover is complete. Since the endpoint for your DB Instance remains the same after a failover, your application can resume database operation without the need for manual administrative intervention.

Enhanced Durability

Multi-AZ deployments for the MySQL, Oracle, and PostgreSQL engines utilize synchronous physical replication to keep data on the standby up-to-date with the primary. Multi-AZ deployments for the SQL Server engine use synchronous logical replication to achieve the same result, employing SQL Server-native Mirroring technology. Both approaches safeguard your data in the event of a DB Instance failure or loss of an Availability Zone.

If a storage volume on your primary fails in a Multi-AZ deployment, Amazon RDS automatically initiates a failover to the up-to-date standby. Compare this to a Single-AZ deployment: in case of a Single-AZ database failure, a user-initiated point-in-time-restore operation will be required. This operation can take several hours to complete, and any data updates that occurred after the latest restorable time (typically within the last five minutes) will not be available.

Amazon Aurora employs a highly durable, SSD-backed virtualized storage layer purpose-built for database workloads. Amazon Aurora automatically replicates your volume six ways, across three Availability Zones. Amazon Aurora storage is fault-tolerant, transparently handling the loss of up to two copies of data without affecting database write availability and up to three copies without affecting read availability. Amazon Aurora storage is also self-healing. Data blocks and disks are continuously scanned for errors and replaced automatically.

Increased Availability

You also benefit from enhanced database availability when running Multi-AZ deployments. If an Availability Zone failure or DB Instance failure occurs, your availability impact is limited to the time automatic failover takes to complete: typically under one minute for Amazon Aurora and one to two minutes for other database engines (see the RDS FAQ for details).

The availability benefits of Multi-AZ deployments also extend to planned maintenance and backups. In the case of system upgrades like OS patching or DB Instance scaling, these operations are applied first on

the standby, prior to the automatic failover. As a result, your availability impact is, again, only the time required for automatic failover to complete.

Unlike Single-AZ deployments, I/O activity is not suspended on your primary during backup for Multi-AZ deployments for the MySQL, Oracle, and PostgreSQL engines, because the backup is taken from the standby. However, note that you may still experience elevated latencies for a few minutes during backups for Multi-AZ deployments.

On instance failure in Amazon Aurora deployments, Amazon RDS uses RDS Multi-AZ technology to automate failover to one of up to 15 Amazon Aurora Replicas you have created in any of three Availability Zones. If no Amazon Aurora Replicas have been provisioned, in the case of a failure, Amazon RDS will attempt to create a new Amazon Aurora DB instance for you automatically.

NEW QUESTION 235

Your application is using an ELB in front of an Auto Scaling group of web/application servers deployed across two AZs and a Multi-AZ RDS Instance for data persistence.

The database CPU is often above 80% usage and 90% of I/O operations on the database are reads. To improve performance you recently added a single-node Memcached ElastiCache Cluster to cache frequent DB query results. In the next weeks the overall workload is expected to grow by 30%.

Do you need to change anything in the architecture to maintain the high availability of the application with the anticipated additional load? Why?

- A. Yes, you should deploy two Memcached ElastiCache Clusters in different AZs because the RDS instance will not be able to handle the load if the cache node fails.
- B. No, if the cache node fails you can always get the same data from the DB without having any availability impact.
- C. No, if the cache node fails the automated ElastiCache node recovery feature will prevent any availability impact.
- D. Yes, you should deploy the Memcached ElastiCache Cluster with two nodes in the same AZ as the RDS DB master instance to handle the load if one cache node fails.

Answer: A

Explanation:

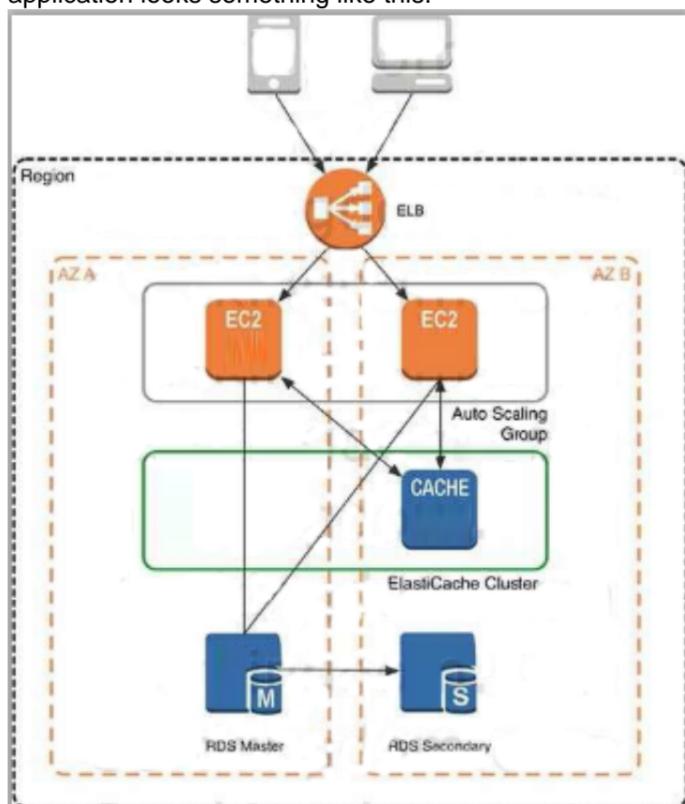
ElastiCache for Memcached

The primary goal of caching is typically to offload reads from your database or other primary data source. In most apps, you have hot spots of data that are regularly queried, but only updated periodically. Think of the front page of a blog or news site, or the top 100 leaderboard in an online game. In this type of case, your app can receive dozens, hundreds, or even thousands of requests for the same data before it's updated again. Having your caching layer handle these queries has several advantages. First, it's considerably cheaper to add an in-memory cache than to scale up to a larger database cluster. Second, an in-memory cache is also easier to scale out, because it's easier to distribute an in-memory cache horizontally than a relational database. Last, a caching layer provides a request buffer in the event of a sudden spike in usage. If your app or game ends up on the front page of Reddit or the App Store, it's not unheard of to see a spike that is 10 to 100 times your normal application load. Even if you autoscale your application instances, a IOx request spike will likely make your database very unhappy.

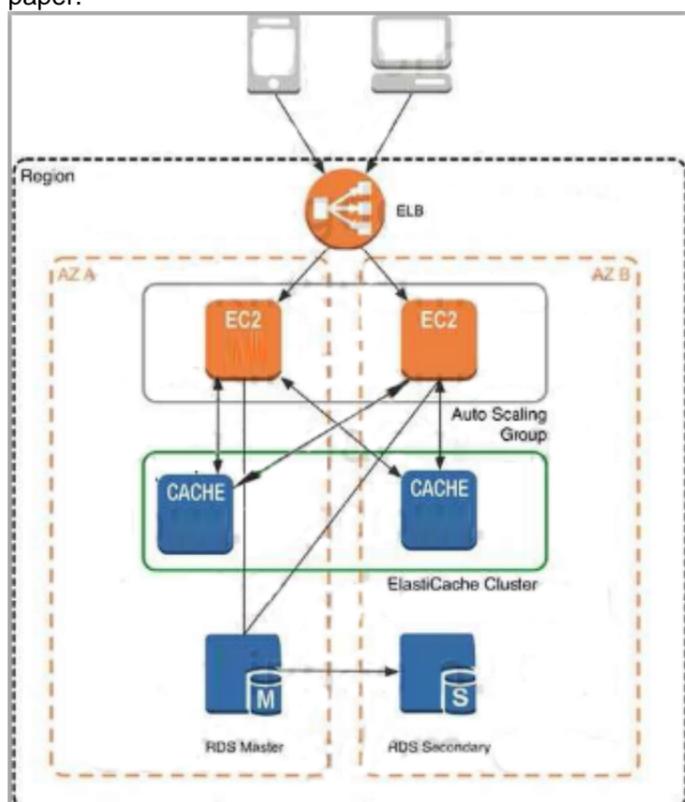
Let's focus on ElastiCache for Memcached first, because it is the best fit for a caching focused solution. We'll revisit Redis later in the paper, and weigh its advantages and disadvantages.

Architecture with ElastiCache for Memcached

When you deploy an ElastiCache Memcached cluster, it sits in your application as a separate tier alongside your database. As mentioned previously, Amazon ElastiCache does not directly communicate with your database tier, or indeed have any particular knowledge of your database. A simplified deployment for a web application looks something like this:



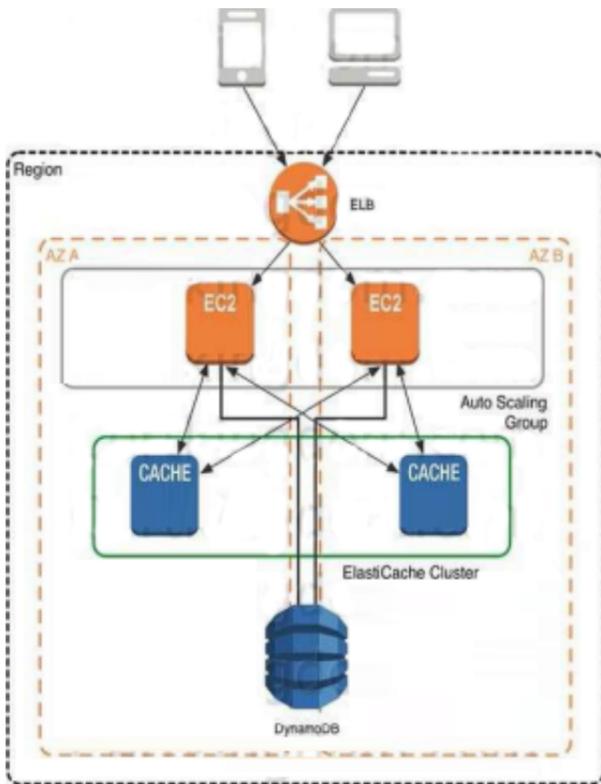
In this architecture diagram, the Amazon EC2 application instances are in an Auto Scaling group, located behind a load balancer using Elastic Load Balancing, which distributes requests among the instances. As requests come into a given EC2 instance, that EC2 instance is responsible for communicating with ElastiCache and the database tier. For development purposes, you can begin with a single ElastiCache node to test your application, and then scale to additional cluster nodes by modifying the ElastiCache cluster. As you add additional cache nodes, the EC2 application instances are able to distribute cache keys across multiple ElastiCache nodes. The most common practice is to use client-side sharding to distribute keys across cache nodes, which we will discuss later in this paper.



When you launch an ElastiCache cluster, you can choose the Availability Zone(s) that the cluster lives in. For best performance, you should configure your cluster to use the same Availability Zones as your application servers. To launch an ElastiCache cluster in a specific Availability Zone, make sure to specify the Preferred Zone(s) option during cache cluster creation. The Availability Zones that you specify will be where ElastiCache will launch your cache nodes. We recommend that you select Spread Nodes Across Zones, which tells ElastiCache to distribute cache nodes across these zones as evenly as possible. This distribution will mitigate the impact of an Availability Zone disruption on your ElastiCache nodes. The trade-off is that some of the requests from your application to ElastiCache will go to a node in a different Availability Zone, meaning latency will be slightly higher.

For more details, refer to Creating a Cache Cluster in the Amazon ElastiCache User Guide.

As mentioned at the outset, ElastiCache can be coupled with a wide variety of databases. Here is an example architecture that uses Amazon DynamoDB instead of Amazon RDS and IVySQL:



This combination of DynamoDB and ElastiCache is very popular with mobile and game companies, because DynamoDB allows for higher write throughput at lower cost than traditional relational databases. In addition, DynamoDB uses a key-value access pattern similar to ElastiCache, which also simplifies the programming model. Instead of using relational SQL for the primary database but then key-value patterns for the cache, both the primary database and cache can be programmed similarly.

In this architecture pattern, DynamoDB remains the source of truth for data, but application reads are offloaded to ElastiCache for a speed boost.

NEW QUESTION 239

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
- B. create a 'Last updated' attribute in your DynamoDB table that would represent the timestamp of the last update and use it as a filter.
- 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. 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.
- E. Send also each write into an SNS queue in the second region; use an auto-scaling group behind the SNS queue to replay the write in the second region.

Answer: A

NEW QUESTION 242

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 two months resulting in significant financial losses. Your CIO is strongly agreeing to move 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
- B. Develop a CloudFormation template which includes your AMI and the required EC2, AutoScaling, and ELB resources to support deploying the application across Multiple- Availability-Zone
- C. Asynchronously replicate transactions from your on-premises database to a database instance in AWS across a secure VPN connection.
- D. Deploy your application on EC2 instances within an Auto Scaling group across multiple availability zone
- E. Asynchronously replicate transactions from your on-premises database to a database instance in AWS across a secure VPN connection.
- F. Create an EBS backed private AMI which includes a fresh install of your application
- G. 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.
- H. Install your application on a compute-optimized EC2 instance capable of supporting the application's average load
- I. Synchronously replicate transactions from your on-premises database to a database instance in AWS across a secure Direct Connect connection.

Answer: A

Explanation:

Overview of Creating Amazon EBS-Backed AMIs

First, launch an instance from an AMI that's similar to the AMI that you'd like to create. You can connect to your instance and customize it. When the instance is configured correctly, ensure data integrity by

stopping the instance before you create an AMI, then create the image. When you create an Amazon EBS-backed AMI, we automatically register it for you.

Amazon EC2 powers down the instance before creating the AMI to ensure that everything on the instance is stopped and in a consistent state during the creation process. If you're confident that your instance is in a consistent state appropriate for AMI creation, you can tell Amazon EC2 not to power down and reboot the instance. Some file systems, such as XFS, can freeze and unfreeze activity, making it safe to create the image without rebooting the instance.

During the AMI-creation process, Amazon EC2 creates snapshots of your instance's root volume and any other EBS volumes attached to your instance. If any volumes attached to the instance are encrypted, the new AMI only launches successfully on instances that support Amazon EBS encryption. For more information, see Amazon EBS Encryption.

Depending on the size of the volumes, it can take several minutes for the AMI-creation process to complete (sometimes up to 24 hours). You may find it more

efficient to create snapshots of your volumes prior to creating your AMI. This way, only small, incremental snapshots need to be created when the AMI is created, and the process completes more quickly (the total time for snapshot creation remains the same). For more information, see [Creating an Amazon EBS Snapshot](#). After the process completes, you have a new AMI and snapshot created from the root volume of the instance. When you launch an instance using the new AMI, we create a new EBS volume for its root volume using the snapshot. Both the AMI and the snapshot incur charges to your account until you delete them. For more information, see [Deregistering Your AMI](#).

If you add instance-store volumes or EBS volumes to your instance in addition to the root device volume, the block device mapping for the new AMI contains information for these volumes, and the block device mappings for instances that you launch from the new AMI automatically contain information for these volumes. The instance-store volumes specified in the block device mapping for the new instance are new and don't contain any data from the instance store volumes of the instance you used to create the AMI. The data on EBS volumes persists. For more information, see [Block Device Mapping](#).

NEW QUESTION 244

Your company hosts a social media site supporting users in multiple countries. You have been asked to provide a highly available design for the application that leverages multiple regions for the most recently accessed content and latency sensitive portions of the website. The most latency sensitive component of the application involves reading user preferences to support web site personalization and ad selection. In addition to running your application in multiple regions, which option will support this application's requirements?

- A. Serve user content from S3. CloudFront and use Route53 latency-based routing between ELBs in each region. Retrieve user preferences from a local DynamoDB table in each region and leverage SQS to capture changes to user preferences with 505 workers for propagating updates to each table.
- B. Use the S3 Copy API to copy recently accessed content to multiple regions and serve user content from S3. CloudFront with dynamic content and an ELB in each region. Retrieve user preferences from an ElastiCache cluster in each region and leverage SNS notifications to propagate user preference changes to a worker node in each region.
- C. Use the S3 Copy API to copy recently accessed content to multiple regions and serve user content from S3. CloudFront and Route53 latency-based routing between ELBs in each region. Retrieve user preferences from a DynamoDB table and leverage SQS to capture changes to user preferences with 505 workers for propagating DynamoDB updates.
- D. Serve user content from S3. CloudFront with dynamic content, and an ELB in each region. Retrieve user preferences from an ElastiCache cluster in each region and leverage Simple Workflow (SWF) to manage the propagation of user preferences from a centralized DB to each ElastiCache cluster.

Answer: A

NEW QUESTION 245

Your company previously configured a heavily used, dynamically routed VPN connection between your on-premises data center and AWS. You recently provisioned a DirectConnect connection and would like to start using the new connection. After configuring DirectConnect settings in the AWS Console, which of the following options will provide the most seamless transition for your users?

- A. Delete your existing VPN connection to avoid routing loops. Configure your DirectConnect router with the appropriate settings and verify network traffic is leveraging DirectConnect.
- B. Configure your DirectConnect router with a higher BGP priority than your VPN router, verify network traffic is leveraging DirectConnect and then delete your existing VPN connection.
- C. Update your VPC route tables to point to the DirectConnect connection. Configure your DirectConnect router with the appropriate settings, verify network traffic is leveraging DirectConnect and then delete the VPN connection.
- D. Configure your DirectConnect router, update your VPC route tables to point to the DirectConnect connection, configure your VPN connection with a higher BGP point, and verify network traffic is leveraging the DirectConnect connection.
- E. And verify network traffic is leveraging the DirectConnect connection.

Answer: D

NEW QUESTION 249

A web company is looking to implement an external payment service into their highly available application deployed in a VPC. Their application EC2 instances are behind a public-facing ELB. Auto scaling is used to add additional instances as traffic increases. Under normal load, the application runs 2 instances in the Auto Scaling group, but at peak it can scale 3x in size. The application instances need to communicate with the payment service over the Internet, which requires whitelisting of all public IP addresses used to communicate with it. A maximum of 4 whitelisting IP addresses are allowed at a time and can be added through an API.

How should they architect their solution?

- A. Route payment requests through two NAT instances setup for High Availability and whitelist the Elastic IP addresses attached to the NAT instances.
- B. Whitelist the VPC Internet Gateway Public IP and route payment requests through the Internet Gateway.
- C. Whitelist the ELB IP addresses and route payment requests from the application servers through the ELB.
- D. Automatically assign public IP addresses to the application instances in the Auto Scaling group and run a script on boot that adds each instance's public IP address to the payment validation whitelist API.

Answer: D

NEW QUESTION 252

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.
- B. Associate the routing table with all VPC subnets.
- C. 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.
- 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.
- E. 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.

Answer: A

NEW QUESTION 255

You have deployed a three-tier web application in a VPC with a CIDR block of 10.0.0.0/28. You initially deploy two web servers, two application servers, two database servers and one NAT instance for a total of seven EC2 instances. The web, application and database servers are deployed across two availability zones (AZs). You also deploy an ELB in front of the two web servers, and use Route53 for DNS. Web traffic gradually increases in the first few days following the deployment, so you attempt to double the number of instances in each tier of the application to handle the new load. Unfortunately, some of these new instances fail to launch.

Which of the following could be the root cause? (Choose 2 answers)

- A. AWS reserves the first and the last private IP address in each subnet's CIDR block so you do not have enough addresses left to launch all of the new EC2 instances
- B. The Internet Gateway (IGW) of your VPC has scaled-up, adding more instances to handle the traffic spike, reducing the number of available private IP addresses for new instance launches
- C. The ELB has scaled-up, adding more instances to handle the traffic spike, reducing the number of available private IP addresses for new instance launches
- D. AWS reserves one IP address in each subnet's CIDR block for Route53 so you do not have enough addresses left to launch all of the new EC2 instances
- E. AWS reserves the first four and the last IP address in each subnet's CIDR block so you do not have enough addresses left to launch all of the new EC2 instances

Answer: CE

NEW QUESTION 259

You are designing Internet connectivity for your VPC. The web servers must be available on the Internet. The application must have a highly available architecture. Which alternatives should you consider? (Choose 2 answers)

- A. Configure a NAT instance in your VPC. Create a default route via the NAT instance and associate it with all subnets. Configure a DNS A record that points to the NAT instance public IP address.
- B. Configure a CloudFront distribution and configure the origin to point to the private IP addresses of your web servers. Configure a Route53 CNAME record to your CloudFront distribution.
- C. Place all your web servers behind ELB. Configure a Route53 CNAME to point to the ELB DNS name.
- D. Assign EIPs to all web servers
- E. Configure a Route53 record set with all EIPs
- F. With health checks and DNS failover.
- G. Configure ELB with an EIP. Place all your web servers behind ELB. Configure a Route53 A record that points to the EIP.

Answer: CD

NEW QUESTION 262

A corporate web application is deployed within an Amazon Virtual Private Cloud (VPC) and is connected to the corporate data center via an IPsec VPN. The application must authenticate against the on-premises LDAP server. After authentication, each logged-in user can only access an Amazon Simple Storage Service (S3) key space specific to that user.

Which two approaches can satisfy these objectives? (Choose 2 answers)

- A. Develop an identity broker that authenticates against IAM Security Token Service to assume a Lambda role in order to get temporary AWS security credentials. The application calls the identity broker to get AWS temporary security credentials with access to the appropriate S3 bucket.
- B. The application authenticates against LDAP and retrieves the name of an IAM role associated with the user.
- C. The application then calls the IAM Security Token Service to assume that IAM role. The application can use the temporary credentials to access the appropriate S3 bucket.
- D. Develop an identity broker that authenticates against LDAP and then calls IAM Security Token Service to get IAM federated user credentials. The application calls the identity broker to get IAM federated user credentials with access to the appropriate S3 bucket.
- E. The application authenticates against LDAP. The application then calls the AWS Identity and Access Management (IAM) Security service to log in to IAM using the LDAP credentials. The application can use the IAM temporary credentials to access the appropriate S3 bucket.
- F. The application authenticates against IAM Security Token Service using the LDAP credentials. The application uses those temporary AWS security credentials to access the appropriate S3 bucket.

Answer: BC

NEW QUESTION 265

You are designing a multi-platform web application for AWS. The application will run on EC2 instances and will be accessed from PCs, tablets and smart phones. Supported accessing platforms are Windows, MacOS, iOS and Android. Separate sticky session and SSL certificate setups are required for different platform types. Which of the following describes the most cost effective and performance efficient architecture setup?

- A. Setup a hybrid architecture to handle session state and SSL certificates on-prem and separate EC2 instance groups running web applications for different platform types running in a VPC
- B. Set up one ELB for all platforms to distribute load among multiple instances under it. Each EC2 instance implements all functionality for a particular platform.
- C. Set up two ELBs. The first ELB handles SSL certificates for all platforms and the second ELB handles session stickiness for all platforms. For each ELB, run separate EC2 instance groups to handle the web application for each platform.
- D. Assign multiple ELBs to an EC2 instance or group of EC2 instances running the common components of the web application, one ELB for each platform type. Session stickiness and SSL termination are done at the ELBs.

Answer: D

NEW QUESTION 270

Your department creates regular analytics reports from your company's log files. All log data is collected in Amazon S3 and processed by daily Amazon Elastic MapReduce (EMR) jobs that generate daily PDF reports and aggregated tables in CSV format for an Amazon Redshift data warehouse.

Your CFO requests that you optimize the cost structure for this system.

Which of the following alternatives will lower costs without compromising average performance of the system or data integrity for the raw data?

- A. Use reduced redundancy storage (RRS) for all data in S3. Use a combination of Spot Instances and Reserved Instances for Amazon EMR jobs
- B. Use Reserved Instances for Amazon Redshift.
- C. Use reduced redundancy storage (RRS) for PDF and .csv data in S3. Add Spot Instances to EMR jobs

- D. Use Spot Instances for Amazon Redshift.
- E. Use reduced redundancy storage (RRS) for PDF and .csv data In Amazon S3. Add Spot Instances to Amazon EMR job
- F. Use Reserved Instances for Amazon Redshift.
- G. Use reduced redundancy storage (RRS) for all data in Amazon S3. Add Spot Instances to Amazon EMR job
- H. Use Reserved Instances for Amazon Redshift.

Answer: C

Explanation:

Using Reduced Redundancy Storage

Amazon S3 stores objects according to their storage class. It assigns the storage class to an object when it is written to Amazon S3. You can assign objects a specific storage class (standard or reduced redundancy) only when you write the objects to an Amazon S3 bucket or when you copy objects that are already stored in Amazon S3. Standard is the default storage class. For information about storage classes, see Object Key and Metadata.

In order to reduce storage costs, you can use reduced redundancy storage for noncritical, reproducible data at lower levels of redundancy than Amazon S3 provides with standard storage. The lower level of redundancy results in less durability and availability, but in many cases, the lower costs can make reduced redundancy storage an acceptable storage solution. For example, it can be a cost effective solution for sharing media content that is durably stored elsewhere. It can also make sense if you are storing thumbnails and other resized images that can be easily reproduced from an original image. Reduced redundancy storage is designed to provide 99.99% durability of objects over a given year.

This durability level corresponds to an average annual expected loss of 0.01% of objects. For example, if you store 10,000 objects using the RRS option, you can, on average, expect to incur an annual loss of a single object per year (0.01% of 10,000 objects).

Note

This annual loss represents an expected average and does not guarantee the loss of less than 0.01% of objects in a given year.

Reduced redundancy storage stores objects on multiple devices across multiple facilities, providing 400 times the durability of a typical disk drive, but it does not replicate objects as many times as Amazon S3 standard storage. In addition, reduced redundancy storage is designed to sustain the loss of data in a single facility.

If an object in reduced redundancy storage has been lost, Amazon S3 will return a 405 error on requests made to that object. Amazon S3 also offers notifications for reduced redundancy storage object loss: you can configure your bucket so that when Amazon S3 detects the loss of an RRS object, a notification will be sent through Amazon Simple Notification Service (Amazon SNS). You can then replace the lost object. To enable notifications, you can use the Amazon S3 console to set the Notifications property of your bucket.

NEW QUESTION 271

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your AWS-Solution-Architect-Associate Exam with Our Prep Materials Via below:

<https://www.certleader.com/AWS-Solution-Architect-Associate-dumps.html>