

Exam Questions AWS-Solution-Architect-Associate

AWS Certified Solutions Architect - Associate

<https://www.2passeasy.com/dumps/AWS-Solution-Architect-Associate/>



NEW QUESTION 1

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 2

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 3

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 4

One of the criteria for a new deployment is that the customer wants to use AWS Storage Gateway. However you are not sure whether you should use gateway-cached volumes or gateway-stored volumes or even what the differences are. Which statement below best describes those differences?

- A. Gateway-cached lets you store your data in Amazon Simple Storage Service (Amazon S3) and retain a copy of frequently accessed data subsets locally
- B. Gateway-stored enables you to configure your on-premises gateway to store all your data locally and then asynchronously back up point-in-time snapshots of this data to Amazon S3.
- C. Gateway-cached is free whilst gateway-stored is not.
- D. Gateway-cached is up to 10 times faster than gateway-stored.
- E. Gateway-stored lets you store your data in Amazon Simple Storage Service (Amazon S3) and retain a copy of frequently accessed data subsets locally
- F. Gateway-cached enables you to configure your on-premises gateway to store all your data locally and then asynchronously back up point-in-time snapshots of this data to Amazon S3.

Answer: A

Explanation: Volume gateways provide cloud-backed storage volumes that you can mount as Internet Small Computer System Interface (iSCSI) devices from your on-premises application servers. The gateway supports the following volume configurations:

Gateway-cached volumes — You store your data in Amazon Simple Storage Service (Amazon S3) and retain a copy of frequently accessed data subsets locally. Gateway-cached volumes offer a substantial cost savings on primary storage and minimize the need to scale your storage on-premises. You also retain low-latency access to your frequently accessed data.

Gateway-stored volumes — If you need low-latency access to your entire data set, you can configure your on-premises gateway to store all your data locally and then asynchronously back up point-in-time snapshots of this data to Amazon S3. This configuration provides durable and inexpensive off-site backups that you can recover to your local data center or Amazon EC2. For example, if you need replacement capacity for disaster recovery, you can recover the backups to Amazon EC2.

Reference: <http://docs.aws.amazon.com/storagegateway/latest/userguide/volume-gateway.html>

NEW QUESTION 5

A user is launching an EC2 instance in the US East region. Which of the below mentioned options is recommended by AWS with respect to the selection of the availability zone?

- A. Always select the AZ while launching an instance

- B. Always select the US-East-1-a zone for HA
- C. Do not select the AZ; instead let AWS select the AZ
- D. The user can never select the availability zone while launching an instance

Answer: C

Explanation: When launching an instance with EC2, AWS recommends not to select the availability zone (AZ). AWS specifies that the default Availability Zone should be accepted. This is because it enables AWS to select the best Availability Zone based on the system health and available capacity. If the user launches additional instances, only then an Availability Zone should be specified. This is to specify the same or different AZ from the running instances.
Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html>

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

You are migrating an internal server on your DC to an EC2 instance with EBS volume. Your server disk usage is around 500GB so you just copied all your data to a 2TB disk to be used with AWS Import/Export. Where will the data be imported once it arrives at Amazon?

- A. to a 2TB EBS volume
- B. to an S3 bucket with 2 objects of 1TB
- C. to an 500GB EBS volume
- D. to an S3 bucket as a 2TB snapshot

Answer: B

Explanation: An import to Amazon EBS will have different results depending on whether the capacity of your storage device is less than or equal to 1 TB or greater than 1 TB. The maximum size of an Amazon EBS snapshot is 1 TB, so if the device image is larger than 1 TB, the image is chunked and stored on Amazon S3. The target location is determined based on the total capacity of the device, not the amount of data on the device.
Reference: <http://docs.aws.amazon.com/AWSImportExport/latest/DG/Concepts.html>

NEW QUESTION 8

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 9

Does DynamoDB support in-place atomic updates?

- A. Yes
- B. No
- C. It does support in-place non-atomic updates
- D. It is not defined

Answer: A

Explanation: DynamoDB supports in-place atomic updates.

Reference:

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/WorkingWithItems.html#WorkingWithItems.AtomicCounters>

NEW QUESTION 10

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

An organization has three separate AWS accounts, one each for development, testing, and production. The organization wants the testing team to have access to certain AWS resources in the production account. How can the organization achieve this?

- A. It is not possible to access resources of one account with another account.
- B. Create the IAM roles with cross account access.
- C. Create the IAM user in a test account, and allow it access to the production environment with the IAM policy.
- D. Create the IAM users with cross account acces

Answer: B

Explanation: An organization has multiple AWS accounts to isolate a development environment from a testing or production environment. At times the users from one account need to access resources in the other account, such as promoting an update from the development environment to the production environment. In this case the IAM role with cross account access will provide a solution. Cross account access lets one account share access to their resources with users in the other AWS accounts.

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

NEW QUESTION 11

You need to import several hundred megabytes of data from a local Oracle database to an Amazon RDS DB instance. What does AWS recommend you use to accomplish this?

- A. Oracle export/import utilities
- B. Oracle SQL Developer
- C. Oracle Data Pump
- D. DBMS_FILE_TRANSFER

Answer: C

Explanation: How you import data into an Amazon RDS DB instance depends on the amount of data you have and the number and variety of database objects in your database.

For example, you can use Oracle SQL Developer to import a simple, 20 MB database; you want to use Oracle Data Pump to import complex databases or databases that are several hundred megabytes or several terabytes in size.

Reference: <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Oracle.Procedural.Importing.html>

NEW QUESTION 13

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 16

You are in the process of creating a Route 53 DNS failover to direct traffic to two EC2 zones. Obviously, if one fails, you would like Route 53 to direct traffic to the other region. Each region has an ELB with some instances being distributed. What is the best way for you to configure the Route 53 health check?

- A. Route 53 doesn't support ELB with an internal health check. You need to create your own Route 53 health check of the ELB
- B. Route 53 natively supports ELB with an internal health chec
- C. Turn "Evaluate target health" off and "Associate with Health Check" on and R53 will use the ELB's internal health check.
- D. Route 53 doesn't support ELB with an internal health chec
- E. You need to associate your resource record set for the ELB with your own health check
- F. Route 53 natively supports ELB with an internal health chec
- G. Turn "Evaluate target health" on and "Associate with Health Check" off and R53 will use the ELB's internal health check.

Answer: D

Explanation: With DNS Failover, Amazon Route 53 can help detect an outage of your website and redirect your end users to alternate locations where your application is operating properly. When you enable this feature, Route 53 uses health checks-regularly making Internet requests to your application's endpoints from multiple locations around the world-to determine whether each endpoint of your application is up or down. To enable DNS Failover for an ELB endpoint, create an Alias record pointing to the ELB and set the "Evaluate Target Health" parameter to true. Route 53 creates and manages the health checks for your ELB automatically. You do not need to create your own Route 53 health check of the ELB. You also do not need to associate your resource record set for the ELB with your own health check, because Route 53 automatically associates it with the health checks that Route 53 manages on your behalf. The ELB health check will also inherit the health of your backend instances behind that ELB.

Reference:

<http://aws.amazon.com/about-aws/whats-new/2013/05/30/amazon-route-53-adds-elb-integration-for-dns-failover/>

NEW QUESTION 17

An Elastic IP address (EIP) is a static IP address designed for dynamic cloud computing. With an EIP, you can mask the failure of an instance or software by rapidly remapping the address to another instance in your account. Your EIP is associated with your AWS account, not a particular EC2 instance, and it remains associated with your account until you choose to explicitly release it. By default how many EIPs is each AWS account limited to on a per region basis?

- A. 1
- B. 5
- C. Unlimited
- D. 10

Answer: B

Explanation: By default, all AWS accounts are limited to 5 Elastic IP addresses per region for each AWS account, because public (IPv4) Internet addresses are a scarce public resource. AWS strongly encourages you to use an EIP primarily for load balancing use cases, and use DNS hostnames for all other inter-node communication.

If you feel your architecture warrants additional EIPs, you would need to complete the Amazon EC2 Elastic IP Address Request Form and give reasons as to your need for additional addresses. Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-ip-addresses-eip.html#using-instance-addressing-limit>

NEW QUESTION 22

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 lost

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 25

You are setting up a VPC and you need to set up a public subnet within that VPC. Which following requirement must be met for this subnet to be considered a public subnet?

- A. Subnet's traffic is not routed to an internet gateway but has its traffic routed to a virtual private gateway.
- B. Subnet's traffic is routed to an internet gateway.
- C. Subnet's traffic is not routed to an internet gateway.
- D. None of these answers can be considered a public subnet

Answer: B

Explanation: A virtual private cloud (VPC) is a virtual network dedicated to your AWS account. It is logically isolated from other virtual networks in the AWS cloud. You can launch your AWS resources, such as Amazon EC2 instances, into your VPC. You can configure your VPC: you can select its IP address range, create subnets, and configure route tables, network gateways, and security settings.

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.

If a subnet's traffic is routed to an internet gateway, the subnet is known as a public subnet.

If a subnet doesn't have a route to the internet gateway, the subnet is known as a private subnet.

If a subnet doesn't have a route to the internet gateway, but has its traffic routed to a virtual private gateway, the subnet is known as a VPN-only subnet.

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

NEW QUESTION 30

You have been asked to build a database warehouse using Amazon Redshift. You know a little about it, including that it is a SQL data warehouse solution, and uses industry standard ODBC and JDBC connections and PostgreSQL drivers. However you are not sure about what sort of storage it uses for database tables. What sort of storage does Amazon Redshift use for database tables?

- A. InnoDB Tables
- B. NDB data storage

- C. Columnar data storage
- D. NDB CLUSTER Storage

Answer: C

Explanation: Amazon Redshift achieves efficient storage and optimum query performance through a combination of massively parallel processing, columnar data storage, and very efficient, targeted data compression encoding schemes.

Columnar storage for database tables is an important factor in optimizing analytic query performance because it drastically reduces the overall disk I/O requirements and reduces the amount of data you need to load from disk.

Reference: http://docs.aws.amazon.com/redshift/latest/dg/c_columnar_storage_disk_mem_mgmt.html

NEW QUESTION 31

You are checking the workload on some of your General Purpose (SSD) and Provisioned IOPS (SSD) volumes and it seems that the I/O latency is higher than you require. You should probably check the to make sure that your application is not trying to drive more IOPS than you have provisioned.

- A. Amount of IOPS that are available
- B. Acknowledgement from the storage subsystem
- C. Average queue length
- D. Time it takes for the I/O operation to complete

Answer: C

Explanation: In EBS workload demand plays an important role in getting the most out of your General Purpose (SSD) and Provisioned IOPS (SSD) volumes. In order for your volumes to deliver the amount of IOPS that are available, they need to have enough I/O requests sent to them. There is a relationship between the demand on the volumes, the amount of IOPS that are available to them, and the latency of the request (the amount of time it takes for the I/O operation to complete).

Latency is the true end-to-end client time of an I/O operation; in other words, when the client sends a IO, how long does it take to get an acknowledgement from the storage subsystem that the IO read or write is complete.

If your I/O latency is higher than you require, check your average queue length to make sure that your application is not trying to drive more IOPS than you have provisioned. You can maintain high IOPS while keeping latency down by maintaining a low average queue length (which is achieved by provisioning more IOPS for your volume).

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

NEW QUESTION 34

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 39

You are building infrastructure for a data warehousing solution and an extra request has come through that there will be a lot of business reporting queries running all the time and you are not sure if your current DB instance will be able to handle it. What would be the best solution for this?

- A. DB Parameter Groups
- B. Read Replicas
- C. Multi-AZ DB Instance deployment
- D. Database Snapshots

Answer: B

Explanation: Read Replicas make it easy to take advantage of MySQL's built-in replication functionality to elastically scale out beyond the capacity constraints of a single DB Instance for read-heavy database workloads. There are a variety of scenarios where deploying one or more Read Replicas for a given source DB Instance may make sense. Common reasons for deploying a Read Replica include:

Scaling beyond the compute or I/O capacity of a single DB Instance for read-heavy database workloads. This excess read traffic can be directed to one or more Read Replicas.

Serving read traffic while the source DB Instance is unavailable. If your source DB Instance cannot take I/O requests (e.g. due to I/O suspension for backups or scheduled maintenance), you can direct read traffic to your Read Replica(s). For this use case, keep in mind that the data on the Read Replica may be "stale" since the source DB Instance is unavailable.

Business reporting or data warehousing scenarios; you may want business reporting queries to run against a Read Replica, rather than your primary, production DB Instance.

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

NEW QUESTION 41

In DynamoDB, could you use IAM to grant access to Amazon DynamoDB resources and API actions?

- A. In DynamoDB there is no need to grant access

- B. Depended to the type of access
- C. No
- D. Yes

Answer: D

Explanation: Amazon DynamoDB integrates with AWS Identity and Access Management (IAM). You can use AWS IAM to grant access to Amazon DynamoDB resources and API actions. To do this, you first write an AWS IAM policy, which is a document that explicitly lists the permissions you want to grant. You then attach that policy to an AWS IAM user or role.

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

NEW QUESTION 44

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 48

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 52

is a fast, filexible, fully managed push messaging service.

- A. Amazon SNS
- B. Amazon SES
- C. Amazon SQS
- D. Amazon FPS

Answer: A

Explanation: Amazon Simple Notification Service (Amazon SNS) is a fast, filexible, 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/?nc1=h_l2_as

NEW QUESTION 55

You have just been given a scope for a new client who has an enormous amount of data(petabytes) that he constantly needs analysed. Currently he is paying a huge amount of money for a data warehousing company to do this for him and is wondering if AWS can provide a cheaper solution. Do you think AWS has a solution for this?

- A. Ye
- B. Amazon SimpleDB
- C. N
- D. Not presently

- E. Ye
- F. Amazon Redshift
- G. Ye
- H. Your choice of relational AMLs on Amazon EC2 and EBS

Answer: C

Explanation: Amazon Redshift is a fast, fully managed, petabyte-scale data warehouse service that makes it simple and cost-effective to efficiently analyze all your data using your existing business intelligence tools. You can start small for just \$0.25 per hour with no commitments or upfront costs and scale to a petabyte or more for \$1,000 per terabyte per year, less than a tenth of most other data warehousing solutions. 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.

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

NEW QUESTION 56

An organization has created an application which is hosted on the AWS EC2 instance. The application stores images to S3 when the end user uploads to it. The organization does not want to store the AWS secure credentials required to access the S3 inside the instance. Which of the below mentioned options is a possible solution to avoid any security threat?

- A. Use the IAM based single sign between the AWS resources and the organization application.
- B. Use the IAM role and assign it to the instance.
- C. Since the application is hosted on EC2, it does not need credentials to access S3.
- D. Use the X.509 certificates instead of the access and the secret access key

Answer: B

Explanation: The AWS IAM role uses temporary security credentials to access AWS services. Once the role is assigned to an instance, it will not need any security credentials to be stored on the instance. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

NEW QUESTION 60

Can resource record sets in a hosted zone have a different domain suffix (for example, www.blog. acme.com and www.acme.ca)?

- A. Yes, it can have for a maximum of three different TLDs.
- B. Yes
- C. Yes, it can have depending on the TLD.
- D. No

Answer: D

Explanation: The resource record sets contained in a hosted zone must share the same suffix. For example, the example.com hosted zone can contain resource record sets for www.example.com and www.aws.example.com subdomains, but it cannot contain resource record sets for a www.example.ca subdomain.

Reference: <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/AboutHostedZones.html>

NEW QUESTION 62

Which of the following AWS CLI commands is syntactically incorrect?

1. \$ aws ec2 describe-instances
2. \$ aws ec2 start-instances --instance-ids i-1348636c
3. \$ aws sns publish --topic-arn arn:aws:sns:us-east-1:546419318123:OperationsError -message "Script Failure"
4. \$ aws sqs receive-message --queue-url https://queue.amazonaws.com/546419318123/Test

- A. 3
- B. 4
- C. 2
- D. 1

Answer: A

Explanation: The following CLI command is missing a hyphen before "-message".

```
aws sns publish --topic-arn arn:aws:sns:us-east-1:546419318123:OperationsError -message "Script Failure"
```

It has been added below in red

```
aws sns publish --topic-arn arn:aws:sns:us-east-1:546419318123:OperationsError ---message "Script Failure"
```

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

NEW QUESTION 64

An online gaming site asked you if you can deploy a database that is a fast, highly scalable NoSQL database service in AWS for a new site that he wants to build. Which database should you recommend?

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

Answer: A

Explanation: Amazon DynamoDB is ideal for database applications that require very low latency and predictable performance at any scale but don't need complex querying capabilities like joins or transactions. Amazon DynamoDB is a fully-managed NoSQL database service that offers high performance, predictable throughput and low cost. It is easy to set up, operate, and scale.

With Amazon DynamoDB, you can start small, specify the throughput and storage you need, and easily scale your capacity requirements on the fly. Amazon DynamoDB automatically partitions data over a number of servers to meet your request capacity. In addition, DynamoDB automatically replicates your data synchronously across multiple Availability Zones within an AWS Region to ensure high-availability and data durability.

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

NEW QUESTION 68

You log in to IAM on your AWS console and notice the following message. "Delete your root access keys." Why do you think IAM is requesting this?

- A. Because the root access keys will expire as soon as you log out.
- B. Because the root access keys expire after 1 week.
- C. Because the root access keys are the same for all users.
- D. Because they provide unrestricted access to your AWS resource

Answer: D

Explanation: In AWS an access key is required in order to sign requests that you make using the command-line interface (CLI), using the AWS SDKs, or using direct API calls. Anyone who has the access key for your root account has unrestricted access to all the resources in your account, including billing information. One of the best ways to protect your account is to not have an access key for your root account. We recommend that unless you must have a root access key (this is very rare), that you do not generate one. Instead, AWS best practice is to create one or more AWS Identity and Access Management (IAM) users, give them the necessary permissions, and use IAM users for everyday interaction with AWS.

Reference:

<http://docs.aws.amazon.com/general/latest/gr/aws-access-keys-best-practices.html#root-password>

NEW QUESTION 69

Your company has been storing a lot of data in Amazon Glacier and has asked for an inventory of what is in there exactly. So you have decided that you need to download a vault inventory. Which of the following statements is incorrect in relation to Vault Operations in Amazon Glacier?

- A. You can use Amazon Simple Notification Service (Amazon SNS) notifications to notify you when the job completes.
- B. A vault inventory refers to the list of archives in a vault.
- C. You can use Amazon Simple Queue Service (Amazon SQS) notifications to notify you when the job completes.
- D. Downloading a vault inventory is an asynchronous operatio

Answer: C

Explanation: Amazon Glacier supports various vault operations.

A vault inventory refers to the list of archives in a vault. For each archive in the list, the inventory provides archive information such as archive ID, creation date, and size. Amazon Glacier updates the vault inventory approximately once a day, starting on the day the first archive is uploaded to the vault. A vault inventory must exist for you to be able to download it.

Downloading a vault inventory is an asynchronous operation. You must first initiate a job to download the inventory. After receiving the job request, Amazon Glacier prepares your inventory for download. After the job completes, you can download the inventory data.

Given the asynchronous nature of the job, you can use Amazon Simple Notification Service (Amazon SNS) notifications to notify you when the job completes. You can specify an Amazon SNS topic for each individual job request or configure your vault to send a notification when specific vault events occur. Amazon Glacier prepares an inventory for each vault periodically, every 24 hours. If there have been no archive additions or deletions to the vault since the last inventory, the inventory date is not updated. When you initiate a job for a vault inventory, Amazon Glacier returns the last inventory it generated, which is a point-in-time snapshot and not real-time data. You might not find it useful to retrieve vault inventory for each archive upload. However, suppose you maintain a database on the client-side associating metadata about the archives you upload to Amazon Glacier. Then, you might find the vault inventory useful to reconcile information in your database with the actual vault inventory.

Reference: <http://docs.aws.amazon.com/amazonglacier/latest/dev/working-with-vaults.html>

NEW QUESTION 70

You are setting up a very complex financial services grid and so far it has 5 Elastic IP (EIP) addresses.

You go to assign another EIP address, but all accounts are limited to 5 Elastic IP addresses per region by default, so you aren't able to. What is the reason for this?

- A. For security reasons.
- B. Hardware restrictions.
- C. Public (IPv4) internet addresses are a scarce resource.
- D. There are only 5 network interfaces per instanc

Answer: C

Explanation: Public (IPv4) internet addresses are a scarce resource. There is only a limited amount of public IP space available, and Amazon EC2 is committed to helping use that space efficiently.

By default, all accounts are limited to 5 Elastic IP addresses per region. If you need more than 5 Elastic IP addresses, AWS asks that you apply for your limit to be raised. They will ask you to think through your use case and help them understand your need for additional addresses.

Reference: http://aws.amazon.com/ec2/faqs/#How_many_instances_can_I_run_in_Amazon_EC2

NEW QUESTION 75

What does Amazon DynamoDB provide?

- A. A predictable and scalable MySQL database
- B. A fast and reliable PL/SQL database cluster

- C. A standalone Cassandra database, managed by Amazon Web Services
- D. A fast, highly scalable managed NoSQL database service

Answer: D

Explanation: Amazon DynamoDB is a managed NoSQL database service offered by Amazon. It automatically manages tasks like scalability for you while it provides high availability and durability for your data, allowing you to concentrate in other aspects of your application.

Reference: check link - https://aws.amazon.com/running_databases/

NEW QUESTION 79

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. What formatting is required for this template?

- A. JSON-formatted document
- B. CSS-formatted document
- C. XML-formatted document
- D. HTML-formatted document

Answer: A

Explanation: You can write an AWS CloudFormation template (a JSON-formatted document) in a text editor or pick an existing template. The template describes the resources you want and their settings. For example, suppose you want to create an Amazon EC2. Your template can declare an instance Amazon EC2 and describe its properties, as shown in the following example:

```
{
  "AWSTemplateFormatVersion" : "2010-09-09",
  "Description" : "A simple Amazon EC2 instance", "Resources" : {
    "MyEC2Instance" : {
      "Type" : "AWS::EC2::Instance", "Properties" : {
        "ImageId" : "ami-2f726546", "InstanceType" : "t1.micro"
      }
    }
  }
}
```

Reference:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-what-is-how-does-it-work.html>

NEW QUESTION 83

Does AWS CloudFormation support Amazon EC2 tagging?

- A. Yes, AWS CloudFormation supports Amazon EC2 tagging
- B. No, CloudFormation doesn't support any tagging
- C. No, it doesn't support Amazon EC2 tagging.
- D. It depends if the Amazon EC2 tagging has been defined in the template

Answer: A

Explanation: In AWS CloudFormation, Amazon EC2 resources that support the tagging feature can also be tagged in an AWS template. The tag values can refer to template parameters, other resource names, resource attribute values (e.g. addresses), or values computed by simple functions (e.g., a concatenated list of strings).

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

NEW QUESTION 84

A user is accessing an EC2 instance on the SSH port for IP 10.20.30.40. Which one is a secure way to configure that the instance can be accessed only from this IP?

- A. In the security group, open port 22 for IP 10.20.30.40
- B. In the security group, open port 22 for IP 10.20.30.40/32
- C. In the security group, open port 22 for IP 10.20.30.40/24
- D. In the security group, open port 22 for IP 10.20.30.40/0

Answer: B

Explanation: In AWS EC2, while configuring a security group, the user needs to specify the IP address in CIDR notation. The CIDR IP range 10.20.30.40/32 says it is for a single IP 10.20.30.40. If the user specifies the IP as 10.20.30.40 only, the security group will not accept and ask it in a CIDR format.

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

NEW QUESTION 86

Which of the following statements is true of creating a launch configuration using an EC2 instance?

- A. The launch configuration can be created only using the Query APIs.
- B. Auto Scaling automatically creates a launch configuration directly from an EC2 instance.
- C. A user should manually create a launch configuration before creating an Auto Scaling group.
- D. The launch configuration should be created manually from the AWS CLI

Answer: B

Explanation: You can create an Auto Scaling group directly from an EC2 instance. When you use this feature, Auto Scaling automatically creates a launch configuration for you as well.

Reference:

<http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/create-lc-with-instanceID.html>

NEW QUESTION 90

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 95

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

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

Answer: A

Explanation: The Multi AZ feature allows the user to achieve High Availability. For Multi AZ, Amazon RDS automatically provisions and maintains a synchronous "standby" replica in a different Availability Zone. Reference: <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Welcome.html>

NEW QUESTION 98

A user has attached 1 EBS volume to a VPC instance. The user wants to achieve the best fault tolerance of data possible. Which of the below mentioned options can help achieve fault tolerance?

- A. Attach one more volume with RAID 1 configuration.
- B. Attach one more volume with RAID 0 configuration.
- C. Connect multiple volumes and stripe them with RAID 6 configuration.
- D. Use the EBS volume as a root device

Answer: A

Explanation: The user can join multiple provisioned IOPS volumes together in a RAID 1 configuration to achieve better fault tolerance. RAID 1 does not provide a write performance improvement; it requires more bandwidth than non-RAID configurations since the data is written simultaneously to multiple volumes.

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

NEW QUESTION 99

A user is aware that a huge download is occurring on his instance. He has already set the Auto Scaling policy to increase the instance count when the network I/O increases beyond a certain limit. How can the user ensure that this temporary event does not result in scaling?

- A. The network I/O are not affected during data download
- B. The policy cannot be set on the network I/O
- C. There is no way the user can stop scaling as it is already configured
- D. Suspend scaling

Answer: D

Explanation: The user may want to stop the automated scaling processes on the Auto Scaling groups either to perform manual operations or during emergency situations. To perform this, the user can suspend one or more scaling processes at any time. Once it is completed, the user can resume all the suspended processes. Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AS_Concepts.html

NEW QUESTION 102

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 106

In Amazon EC2, if your EBS volume stays in the detaching state, you can force the detachment by clicking .

- A. Force Detach
- B. Detach Instance
- C. AttachVolume
- D. AttachInstance

Answer: A

Explanation: If your volume stays in the detaching state, you can force the detachment by clicking Force Detach. Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/ebs-detaching-volume.html>

NEW QUESTION 107

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 108

The common use cases for DynamoDB Fine-Grained Access Control (FGAC) are cases in which the end user wants .

- A. to change the hash keys of the table directly
- B. to check if an IAM policy requires the hash keys of the tables directly
- C. to read or modify any codecommit key of the table directly, without a middle-tier service
- D. to read or modify the table directly, without a middle-tier service

Answer: D

Explanation: FGAC can benefit any application that tracks information in a DynamoDB table, where the end user (or application client acting on behalf of an end user) wants to read or modify the table directly, without a middle-tier service. For instance, a developer of a mobile app named Acme can use FGAC to track the top score of every Acme user in a DynamoDB table. FGAC allows the application client to modify only the top score for the user that is currently running the application.

Reference: http://aws.amazon.com/dynamodb/faqs/#security_anchor

NEW QUESTION 110

A user comes to you and wants access to Amazon CloudWatch but only wants to monitor a specific LoadBalancer. Is it possible to give him access to a specific set of instances or a specific LoadBalancer?

- A. No because you can't use IAM to control access to CloudWatch data for specific resources.
- B. Ye
- C. You can use IAM to control access to CloudWatch data for specific resources.
- D. No because you need to be Sysadmin to access CloudWatch data.
- E. Ye
- F. Any user can see all CloudWatch data and needs no access right

Answer: A

Explanation: Amazon CloudWatch integrates with AWS Identity and Access Management (IAM) so that you can specify which CloudWatch actions a user in your AWS Account can perform. For example, you could create an IAM policy that gives only certain users in your organization permission to use GetMetricStatistics. They could then use the action to retrieve data about your cloud resources.

You can't use IAM to control access to CloudWatch data for specific resources. For example, you can't give a user access to CloudWatch data for only a specific set of instances or a specific LoadBalancer. Permissions granted using IAM cover all the cloud resources you use with CloudWatch. In addition, you can't use IAM roles with the Amazon CloudWatch command line tools.

Using Amazon CloudWatch with IAM doesn't change how you use CloudWatch. There are no changes to CloudWatch actions, and no new CloudWatch actions related to users and access control.

Reference: <http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/UsingIAM.html>

NEW QUESTION 115

You have multiple VPN connections and want to provide secure communication between sites using the AWS VPN CloudHub. Which statement is the most accurate in describing what you must do to set this up correctly?

- A. Create a virtual private gateway with multiple customer gateways, each with unique Border Gateway Protocol (BGP) Autonomous System Numbers (ASNs)
- B. Create a virtual private gateway with multiple customer gateways, each with a unique set of keys
- C. Create a virtual public gateway with multiple customer gateways, each with a unique Private subnet
- D. Create a virtual private gateway with multiple customer gateways, each with unique subnet id

Answer: A

Explanation: If you have multiple VPN connections, you can provide secure communication between sites using the AWS VPN CloudHub. The VPN CloudHub operates on a simple hub-and-spoke model that you can use with or without a VPC. This design is suitable for customers with multiple branch offices and existing Internet connections who'd like to implement a convenient, potentially low-cost hub-and-spoke model for primary or backup connectivity between these remote offices.

To use the AWS VPN CloudHub, you must create a virtual private gateway with multiple customer gateways, each with unique Border Gateway Protocol (BGP) Autonomous System Numbers (ASNs). Customer gateways advertise the appropriate routes (BGP prefixes) over their VPN connections. These routing advertisements are received and re-advertised to each BGP peer, enabling each site to send data to and receive data from the other sites. The routes for each spoke must have unique ASNs and the sites must not have overlapping IP ranges. Each site can also send and receive data from the VPC as if they were using a standard VPN connection.

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

NEW QUESTION 117

While creating an Amazon RDS DB, your first task is to set up a DB that controls which IP address or EC2 instance can access your DB Instance.

- A. security token pool
- B. security token
- C. security pool
- D. security group

Answer: D

Explanation: While creating an Amazon RDS DB, your first task is to set up a DB Security Group that controls what IP addresses or EC2 instances have access to your DB Instance.

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

NEW QUESTION 119

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 121

You are very concerned about security on your network because you have multiple programmers testing APIs and SDKs and you have no idea what is happening. You think CloudTrail may help but are not sure what it does. Which of the following statements best describes the AWS service CloudTrail?

- A. With AWS CloudTrail you can get a history of AWS API calls and related events for your account.
- B. With AWS CloudTrail you can get a history of IAM users for your account.
- C. With AWS CloudTrail you can get a history of S3 logfiles for your account.
- D. With AWS CloudTrail you can get a history of CloudFormation JSON scripts used for your account.

Answer: A

Explanation: With AWS CloudTrail, you can get a history of AWS API calls for your account, including API calls made via the AWS Management Console, the AWS SDKs, the command line tools, and higher-level AWS services. You can also identify which users and accounts called AWS APIs for services that support CloudTrail, the source IP address the calls were made from, and when the calls occurred.

You can identify which users and accounts called AWS for services that support CloudTrail, the source IP address the calls were made from, and when the calls occurred. You can integrate CloudTrail into applications using the API, automate trail creation for your organization, check the status of your trails, and control how administrators turn CloudTrail logging on and off.

Reference: http://docs.aws.amazon.com/awscloudtrail/latest/userguide/what_is_cloud_trail_top_level.html

NEW QUESTION 126

Regarding Amazon Route 53, if your application is running on Amazon EC2 instances in two or more Amazon EC2 regions and if you have more than one Amazon EC2 instance in one or more regions, you can use to route traffic to the correct region and then use to route traffic to instances within the region, based on probabilities that you specify.

- A. weighted-based routing; alias resource record sets
- B. latency-based routing; weighted resource record sets
- C. weighted-based routing; weighted resource record sets
- D. latency-based routing; alias resource record sets

Answer: B

Explanation: Regarding Amazon Route 53, if your application is running on Amazon EC2 instances in two or more Amazon EC2 regions, and if you have more than one Amazon EC2 instance in one or more regions, you can use latency-based routing to route traffic to the correct region and then use weighted resource record sets to route traffic to instances within the region based on weights that you specify.

Reference: <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/Tutorials.html>

NEW QUESTION 130

You receive a bill from AWS but are confused because you see you are incurring different costs for the exact same storage size in different regions on Amazon S3. You ask AWS why this is so. What response would you expect to receive from AWS?

- A. We charge less in different time zones.
- B. We charge less where our costs are less.
- C. This will balance out next bill.
- D. It must be a mistake.

Answer: B

Explanation: Amazon S3 is storage for the internet. It's a simple storage service that offers software developers a highly-scalable, reliable, and low-latency data storage infrastructure at very low costs.

AWS charges less where their costs are less.

For example, their costs are lower in the US Standard Region than in the US West (Northern California) Region.

Reference: <https://aws.amazon.com/s3/faqs/>

NEW QUESTION 132

You receive the following request from a client to quickly deploy a static website for them, specifically on AWS. The requirements are low-cost, reliable, online storage, and a reliable and cost-effective way to route customers to the website, as well as a way to deliver content with low latency and high data transfer speeds so that visitors to his website don't experience unnecessary delays. What do you think would be the minimum AWS services that could fulfill the client's request?

- A. Amazon Route 53, Amazon CloudFront and Amazon VPC.
- B. Amazon S3, Amazon Route 53 and Amazon RDS
- C. Amazon S3, Amazon Route 53 and Amazon CloudFront
- D. Amazon S3 and Amazon Route 53.

Answer: C

Explanation: You can easily and inexpensively use AWS to host a website that uses client-side technologies (such as HTML, CSS, and JavaScript) and does not require server-side technologies (such as PHP and ASP.NET). This type of site is called a static website, and is used to display content that does not change frequently. Before you create and deploy a static website, you must plan your architecture to ensure that it meets your requirements. Amazon S3, Amazon Route 53, and Amazon CloudFront would be required in this instance.

Reference: <http://docs.aws.amazon.com/gettingstarted/latest/swf/website-hosting-intro.html>

NEW QUESTION 133

What is the default maximum number of Access Keys per user?

- A. 10
- B. 15
- C. 2
- D. 20

Answer: C

Explanation: The default maximum number of Access Keys per user is 2.

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

NEW QUESTION 134

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 137

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 141

A user is hosting a website in the US West-1 region. The website has the highest client base from the Asia-Pacific (Singapore / Japan) region. The application is accessing data from S3 before serving it to client. Which of the below mentioned regions gives a better performance for S3 objects?

- A. Japan
- B. Singapore
- C. US East
- D. US West-1

Answer: D

Explanation: Access to Amazon S3 from within Amazon EC2 in the same region is fast. In this aspect, though the client base is Singapore, the application is being hosted in the US West-1 region. Thus, it is recommended that S3 objects be stored in the US-West-1 region.
Reference: http://media.amazonwebservices.com/AWS_Storage_Options.pdf

NEW QUESTION 142

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 145

You have three Amazon EC2 instances with Elastic IP addresses in the US East (Virginia) region, and you want to distribute requests across all three IPs evenly for users for whom US East (Virginia) is the appropriate region. How many EC2 instances would be sufficient to distribute requests in other regions?

- A. 3
- B. 9
- C. 2
- D. 1

Answer: D

Explanation: If your application is running on Amazon EC2 instances in two or more Amazon EC2 regions, and if you have more than one Amazon EC2 instance in one or more regions, you can use latency-based routing to route traffic to the correct region and then use weighted resource record sets to route traffic to

instances within the region based on weights that you specify.

For example, suppose you have three Amazon EC2 instances with Elastic IP addresses in the US East (Virginia) region and you want to distribute requests across all three IPs evenly for users for whom US East (Virginia) is the appropriate region. Just one Amazon EC2 instance is sufficient in the other regions, although you can apply the same technique to many regions at once.

Reference: <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/Tutorials.html>

NEW QUESTION 150

In Amazon EC2, what is the limit of Reserved Instances per Availability Zone each month?

- A. 5
- B. 20
- C. 50
- D. 10

Answer: B

Explanation: There are 20 Reserved Instances per Availability Zone in each month.

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

NEW QUESTION 153

The AWS CloudHSM service defines a resource known as a high-availability (HA) , which is a virtual partition that represents a group of partitions, typically distributed between several physical HSMs for high-availability.

- A. proxy group
- B. partition group
- C. functional group
- D. relational group

Answer: B

Explanation: The AWS CloudHSM service defines a resource known as a high-availability (HA) partition group, which is a virtual partition that represents a group of partitions, typically distributed between several physical HSMs for high-availability.

Reference: <http://docs.aws.amazon.com/cloudhsm/latest/userguide/configuring-ha.html>

NEW QUESTION 158

Is it possible to get a history of all EC2 API calls made on your account for security analysis and operational troubleshooting purposes?

- A. Yes, by default, the history of your API calls is logged.
- B. Yes, you should turn on the CloudTrail in the AWS console.
- C. No, you can only get a history of VPC API calls.
- D. No, you cannot store history of EC2 API calls on Amazon.

Answer: B

Explanation: To get a history of all EC2 API calls (including VPC and EBS) made on your account, you simply turn on CloudTrail in the AWS Management Console.

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

NEW QUESTION 160

You have just set up your first Elastic Load Balancer (ELB) but it does not seem to be configured properly. You discover that before you start using ELB, you have to configure the listeners for your load balancer. Which protocols does ELB use to support the load balancing of applications?

- A. HTTP and HTTPS
- B. HTTP, HTTPS , TCP, SSL and SSH
- C. HTTP, HTTPS , TCP, and SSL
- D. HTTP, HTTPS , TCP, SSL and SFTP

Answer: C

Explanation: Before you start using Elastic Load Balancing(ELB), you have to configure the listeners for your load balancer. A listener is a process that listens for connection requests. It is configured with a protocol and a port number for front-end (client to load balancer) and back-end (load balancer to back-end instance) connections.

Elastic Load Balancing supports the load balancing of applications using HTTP, HTTPS (secure HTTP), TCP, and SSL (secure TCP) protocols. The HTTPS uses the SSL protocol to establish secure connections over the HTTP layer. You can also use SSL protocol to establish secure connections over the TCP layer.

The acceptable ports for both HTTPS/SSL and HTTP/TCP connections are 25, 80, 443, 465, 587, and 1024-65535.

Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/elb-listener-config.html>

NEW QUESTION 161

After setting up some EC2 instances you now need to set up a monitoring solution to keep track of these instances and to send you an email when the CPU hits a certain threshold. Which statement below best describes what thresholds you can set to trigger a CloudWatch Alarm?

- A. Set a target value and choose whether the alarm will trigger when the value is greater than (>), greater than or equal to (>=), less than (<), or less than or equal to (<=) that value.
- B. Thresholds need to be set in IAM not CloudWatch
- C. Only default thresholds can be set you can't choose your own thresholds.
- D. Set a target value and choose whether the alarm will trigger when the value hits this threshold

Answer: A

Explanation: Amazon CloudWatch is a monitoring service for AWS cloud resources and the applications you run on AWS. You can use Amazon CloudWatch to collect and track metrics, collect and monitor log files, and set alarms.

When you create an alarm, you first choose the Amazon CloudWatch metric you want it to monitor. Next, you choose the evaluation period (e.g., five minutes or one hour) and a statistical value to measure (e.g., Average or Maximum).

To set a threshold, set a target value and choose whether the alarm will trigger when the value is greater than (>), greater than or equal to (>=), less than (<), or less than or equal to (<=) that value.

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

NEW QUESTION 166

When does the billing of an Amazon EC2 system begin?

- A. It starts when the Status column for your distribution changes from Creating to Deployed.
- B. It starts as soon as you click the create instance option on the main EC2 console.
- C. It starts when your instance reaches 720 instance hours.
- D. It starts when Amazon EC2 initiates the boot sequence of an AM instance

Answer: D

Explanation: Billing commences when Amazon EC2 initiates the boot sequence of an AM instance. Billing ends when the instance terminates, which could occur through a web services command, by running "shutdown -h", or through instance failure. When you stop an instance, Amazon shuts it down but doesn't charge hourly usage for a stopped instance, or data transfer fees, but charges for the storage for any Amazon EBS volumes.

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

NEW QUESTION 168

What is the data model of DynamoDB?

- A. Since DynamoDB is schema-less, there is no data model.
- B. "Items", with Keys and one or more Attribute; and "Attribute", with Name and Value.
- C. "Table", a collection of Items; "Items", with Keys and one or more Attribute; and "Attribute", with Name and Value.
- D. "Database", which is a set of "Tables", which is a set of "Items", which is a set of "Attributes".

Answer: C

Explanation: The data model of DynamoDB is: "Table", a collection of Items; "Items", with Keys and one or more Attribute; "Attribute", with Name and Value.

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

NEW QUESTION 171

What happens to Amazon EBS root device volumes, by default, when an instance terminates?

- A. Amazon EBS root device volumes are moved to IAM.
- B. Amazon EBS root device volumes are copied into Amazon RDS.
- C. Amazon EBS root device volumes are automatically deleted.
- D. Amazon EBS root device volumes remain in the database until you delete the

Answer: C

Explanation: By default, Amazon EBS root device volumes are automatically deleted when the instance terminates. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/terminating-instances.html>

NEW QUESTION 176

Which of the following would you use to list your AWS Import/Export jobs?

- A. Amazon RDS
- B. AWS Import/Export Web Service Tool
- C. Amazon S3 REST API
- D. AWS Elastic Beanstalk

Answer: C

Explanation: You can list AWS Import/Export jobs with the ListJobs command using the command line client or REST API.

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

NEW QUESTION 178

A gaming company comes to you and asks you to build them infrastructure for their site. They are not sure how big they will be as with all start ups they have limited money and big ideas. What they do tell you is that if the game becomes successful, like one of their previous games, it may rapidly grow to millions of users and generate tens (or even hundreds) of thousands of writes and reads per second. After considering all of this, you decide that they need a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. Which of the following databases do you think would best fit their needs?

- A. Amazon DynamoDB
- B. Amazon Redshift
- C. Any non-relational database.
- D. Amazon SimpleDB

Answer: A

Explanation: Amazon DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. Amazon DynamoDB enables customers to offload the administrative burdens of operating and scaling distributed databases to AWS, so they don't have to worry about hardware provisioning, setup and configuration, replication, software patching, or cluster scaling. Today's web-based applications generate and consume massive amounts of data. For example, an online game might start out with only a few thousand users and a light database workload consisting of 10 writes per second and 50 reads per second. However, if the game becomes successful, it may rapidly grow to millions of users and generate tens (or even hundreds) of thousands of writes and reads per second. It may also create terabytes or more of data per day. Developing your applications against Amazon DynamoDB enables you to start small and simply dial-up your request capacity for a table as your requirements scale, without incurring downtime. You pay highly cost-efficient rates for the request capacity you provision, and let Amazon DynamoDB do the work over partitioning your data and traffic over sufficient server capacity to meet your needs. Amazon DynamoDB does the database management and administration, and you simply store and request your data. Automatic replication and failover provides built-in fault tolerance, high availability, and data durability. Amazon DynamoDB gives you the peace of mind that your database is fully managed and can grow with your application requirements.
Reference: <http://aws.amazon.com/dynamodb/faqs/>

NEW QUESTION 180

A favored client needs you to quickly deploy a database that is a relational database service with minimal administration as he wants to spend the least amount of time administering it. Which database would be the best option?

- A. Amazon SimpleDB
- B. Your choice of relational AMs on Amazon EC2 and EBS.
- C. Amazon RDS
- D. Amazon Redshift

Answer: C

Explanation: Amazon Relational Database Service (Amazon RDS) is a web service that makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity while managing time-consuming database administration tasks, freeing you up to focus on your applications and business. 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: https://aws.amazon.com/running_databases/#rds_anchor

NEW QUESTION 182

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 start to look into security groups first. Which statement below is incorrect in relation to security groups?

- A. Are stateful: Return traffic is automatically allowed, regardless of any rules.
- B. Evaluate all rules before deciding whether to allow traffic.
- C. Support allow rules and deny rules.
- D. Operate at the instance level (first layer of defense).

Answer: C

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 and supports allow rules only.
Network access control lists (ACLs)—Act as a firewall for associated subnets, controlling both inbound and outbound traffic at the subnet level and supports allow rules and deny rules.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Security.html

NEW QUESTION 184

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 185

Amazon Elastic Load Balancing is used to manage traffic on a fleet of Amazon EC2 instances, distributing traffic to instances across all availability zones within a region. Elastic Load Balancing has all the advantages of an on-premises load balancer, plus several security benefits.

Which of the following is not an advantage of ELB over an on-premise load balancer?

- A. ELB uses a four-tier, key-based architecture for encryption.
- B. ELB offers clients a single point of contact, and can also serve as the first line of defense against attacks on your network.
- C. ELB takes over the encryption and decryption work from the Amazon EC2 instances and manages it centrally on the load balancer.
- D. ELB supports end-to-end traffic encryption using TLS (previously SSL) on those networks that use secure HTTP (HTTPS) connections.

Answer: A

Explanation: Amazon Elastic Load Balancing is used to manage traffic on a fleet of Amazon EC2 instances, distributing traffic to instances across all availability zones within a region. Elastic Load Balancing has all the advantages of an on-premises load balancer, plus several security benefits:

Takes over the encryption and decryption work from the Amazon EC2 instances and manages it centrally on the load balancer

Offers clients a single point of contact, and can also serve as the first line of defense against attacks on your network

When used in an Amazon VPC, supports creation and management of security groups associated with your Elastic Load Balancing to provide additional networking and security options

Supports end-to-end traffic encryption using TLS (previously SSL) on those networks that use secure HTTP (HTTPS) connections. When TLS is used, the TLS server certificate used to terminate client connections can be managed centrally on the load balancer, rather than on every individual instance. Reference:

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

NEW QUESTION 189

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 193

An organization has a statutory requirement to protect the data at rest for the S3 objects. Which of the below mentioned options need not be enabled by the organization to achieve data security?

- A. MFA delete for S3 objects
- B. Client side encryption
- C. Bucket versioning
- D. Data replication

Answer: D

Explanation: AWS S3 provides multiple options to achieve the protection of data at REST. The options include Permission (Policy), Encryption (Client and Server Side), Bucket Versioning and MFA based delete. The user can enable any of these options to achieve data protection. Data replication is an internal facility by AWS where S3 replicates each object across all the Availability Zones and the organization need not enable it in this case.

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

NEW QUESTION 195

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 197

What is the time period with which metric data is sent to CloudWatch when detailed monitoring is enabled on an Amazon EC2 instance?

- A. 15 minutes
- B. 5 minutes
- C. 1 minute
- D. 45 seconds

Answer: C

Explanation: By default, Amazon EC2 metric data is automatically sent to CloudWatch in 5-minute periods. However, you can, enable detailed monitoring on an Amazon EC2 instance, which sends data to CloudWatch in 1-minute periods

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

NEW QUESTION 200

Which of the following features are provided by Amazon EC2?

- A. Exadata Database Machine, Optimized Storage Management, Flashback Technology, and Data Warehousing
- B. Instances, Amazon Machine Images (AMIs), Key Pairs, Amazon EBS Volumes, Firewall, Elastic IP address, Tags, and Virtual Private Clouds (VPCs)
- C. Real Application Clusters (RAC), ElastiCache Machine Images (EMIs), Data Warehousing, Flashback Technology, Dynamic IP address
- D. Exadata Database Machine, Real Application Clusters (RAC), Data Guard, Table and Index Partitioning, and Data Pump Compression

Answer: B

Explanation: Amazon EC2 provides the following features:

- Virtual computing environments, known as instances;
- Pre-configured templates for your instances, known as Amazon Machine Images (AMIs), that package the bits you need for your server (including the operating system and additional software)
- Various configurations of CPU, memory, storage, and networking capacity for your instances, known as instance types
- Secure login information for your instances using key pairs (AWS stores the public key, and you store the private key in a secure place)
- Storage volumes for temporary data that's deleted when you stop or terminate your instance, known as instance store volumes
- Persistent storage volumes for your data using Amazon Elastic Block Store (Amazon EBS), known as Amazon EBS volumes
- Multiple physical locations for your resources, such as instances and Amazon EBS volumes, known as regions and Availability Zones
- A firewall that enables you to specify the protocols, ports, and source IP ranges that can reach your instances using security groups
- Static IP addresses for dynamic cloud computing, known as Elastic IP addresses
- Metadata, known as tags, that you can create and assign to your Amazon EC2 resources
- Virtual networks you can create that are logically isolated from the rest of the AWS cloud, and that you can optionally connect to your own network, known as virtual private clouds (VPCs).

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

NEW QUESTION 202

In Amazon Elastic Compute Cloud, which of the following is used for communication between instances in the same network (EC2-Classic or a VPC)?

- A. Private IP addresses
- B. Elastic IP addresses
- C. Static IP addresses
- D. Public IP addresses

Answer: A

Explanation: A private IP address is an IP address that's not reachable over the Internet. You can use private IP addresses for communication between instances in the same network (EC2-Classic or a VPC). Reference:

<http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/using-instance-addressing.html>

NEW QUESTION 207

A friend tells you he is being charged \$100 a month to host his WordPress website, and you tell him you can move it to AWS for him and he will only pay a fraction of that, which makes him very happy. He then tells you he is being charged \$50 a month for the domain, which is registered with the same people that set it up, and he asks if it's possible to move that to AWS as well. You tell him you aren't sure, but will look into it. Which of the following statements is true in regards to transferring domain names to AWS?

- A. You can't transfer existing domains to AWS.
- B. You can transfer existing domains into Amazon Route 53's management.
- C. You can transfer existing domains via AWS Direct Connect.
- D. You can transfer existing domains via AWS Import/Export

Answer: B

Explanation: With Amazon Route 53, you can create and manage your public DNS records with the AWS Management Console or with an easy-to-use API. If you need a domain name, you can find an available name and register it using Amazon Route 53. You can also transfer existing domains into Amazon Route 53's

management.

Reference: <http://aws.amazon.com/route53/>

NEW QUESTION 212

Are penetration tests allowed as long as they are limited to the customer's instances?

- A. Yes, they are allowed but only for selected regions.
- B. No, they are never allowed.
- C. Yes, they are allowed without any permission.
- D. Yes, they are allowed but only with approval.

Answer: D

Explanation: Penetration tests are allowed after obtaining permission from AWS to perform them. Reference: <http://aws.amazon.com/security/penetration-testing/>

NEW QUESTION 215

A user has created an ELB with the availability zone US-East-1A. The user wants to add more zones to ELB to achieve High Availability. How can the user add more zones to the existing ELB?

- A. The user should stop the ELB and add zones and instances as required
- B. The only option is to launch instances in different zones and add to ELB
- C. It is not possible to add more zones to the existing ELB
- D. The user can add zones on the fly from the AWS console

Answer: D

Explanation: The user has created an Elastic Load Balancer with the availability zone and wants to add more zones to the existing ELB. The user can do so in two ways:

From the console or CLI, add new zones to ELB;

Launch instances in a separate AZ and add instances to the existing ELB. Reference:

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

NEW QUESTION 216

What happens to data on an ephemeral volume of an EBS-backed EC2 instance if it is terminated or if it fails?

- A. Data is automatically copied to another volume.
- B. The volume snapshot is saved in S3.
- C. Data persists.
- D. Data is delete

Answer: D

Explanation: Any data on the instance store volumes persists as long as the instance is running, but this data is deleted when the instance is terminated or if it fails (such as if an underlying drive has issues). After an instance store-backed instance fails or terminates, it cannot be restored.

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

NEW QUESTION 221

In AWS CloudHSM, in addition to the AWS recommendation that you use two or more HSM appliances in a high-availability configuration to prevent the loss of keys and data, you can also perform a remote backup/restore of a Luna SA partition if you have purchased a:

- A. Luna Restore HSNI.
- B. Luna Backup HSM.
- C. Luna HSNI.
- D. Luna SA HSM.

Answer: B

Explanation: In AWS CloudHSM, you can perform a remote backup/restore of a Luna SA partition if you have purchased a Luna Backup HSM.

Reference: <http://docs.aws.amazon.com/cloudhsm/latest/userguide/cloud-hsm-backup-restore.html>

NEW QUESTION 222

A user has launched a large EBS backed EC2 instance in the US-East-1a region. The user wants to achieve Disaster Recovery (DR) for that instance by creating another small instance in Europe. How can the user achieve DR?

- A. Copy the instance from the US East region to the EU region
- B. Use the "Launch more like this" option to copy the instance from one region to another
- C. Copy the running instance using the "Instance Copy" command to the EU region
- D. Create an AMI of the instance and copy the AMI to the EU regio
- E. Then launch the instance from the EU AMI

Answer: D

Explanation: To launch an EC2 instance it is required to have an AMI in that region. If the AMI is not available in that region, then create a new AMI or use the copy command to copy the AMI from one region to the other region.

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

NEW QUESTION 226

AWS Identity and Access Management is a web service that enables Amazon Web Services (AWS) customers to manage users and user permissions in AWS. In addition to supporting IAM user policies, some services support resource-based permissions. Which of the following services are supported by resource-based permissions?

- A. Amazon SNS, and Amazon SQS and AWS Direct Connect.
- B. Amazon S3 and Amazon SQS and Amazon ElastiCache.
- C. Amazon S3, Amazon SNS, Amazon SQS, Amazon Glacier and Amazon EBS.
- D. Amazon Glacier, Amazon SNS, and Amazon CloudWatch

Answer: C

Explanation: In addition to supporting IAM user policies, some services support resource-based permissions, which let you attach policies to the service's resources instead of to IAM users or groups. Resource-based permissions are supported by Amazon S3, Amazon SNS, Amazon SQS, Amazon Glacier and Amazon EBS.

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

NEW QUESTION 231

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 234

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.htm>

NEW QUESTION 235

Your company has multiple IT departments, each with their own VPC. Some VPCs are located within the same AWS account, and others in a different AWS account. You want to peer together all VPCs to enable the IT departments to have full access to each others' resources. There are certain limitations placed on VPC peering. Which of the following statements is incorrect in relation to VPC peering?

- A. Private DNS values cannot be resolved between instances in peered VPCs.
- B. You can have up to 3 VPC peering connections between the same two VPCs at the same time.
- C. You cannot create a VPC peering connection between VPCs in different regions.
- D. You have a limit on the number active and pending VPC peering connections that you can have per VPC.

Answer: B

Explanation: To create a VPC peering connection with another VPC, you need to be aware of the following limitations and rules:

You cannot create a VPC peering connection between VPCs that have matching or overlapping CIDR blocks.

You cannot create a VPC peering connection between VPCs in different regions.

You have a limit on the number active and pending VPC peering connections that you can have per VPC. VPC peering does not support transitive peering relationships; in a VPC peering connection, your VPC will not have access to any other VPCs that the peer VPC may be peered with. This includes VPC peering connections that are established entirely within your own AWS account.

You cannot have more than one VPC peering connection between the same two VPCs at the same time. The Maximum Transmission Unit (MTU) across a VPC peering connection is 1500 bytes.

A placement group can span peered VPCs; however, you will not get full-bisection bandwidth between instances in peered VPCs.

Unicast reverse path forwarding in VPC peering connections is not supported.

You cannot reference a security group from the peer VPC as a source or destination for ingress or egress rules in your security group. Instead, reference CIDR blocks of the peer VPC as the source or destination of your security group's ingress or egress rules.

Private DNS values cannot be resolved between instances in peered VPCs. Reference:

<http://docs.aws.amazon.com/AmazonVPC/latest/PeeringGuide/vpc-peering-overview.html#vpc-peering-limitations>

NEW QUESTION 240

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 244

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 249

In the most recent company meeting, your CEO focused on the fact that everyone in the organization needs to make sure that all of the infrastructure that is built is truly scalable. Which of the following statements is incorrect in reference to scalable architecture?

- A. A scalable service is capable of handling heterogeneity.
- B. A scalable service is resilient.
- C. A scalable architecture won't be cost effective as it grows.
- D. Increasing resources results in a proportional increase in performance

Answer: C

Explanation: In AWS it is critical to build a scalable architecture in order to take advantage of a scalable infrastructure. The cloud is designed to provide conceptually infinite scalability. However, you cannot leverage all that scalability in infrastructure if your architecture is not scalable. Both have to work together. You will have to identify the monolithic components and bottlenecks in your architecture, identify the areas where you cannot leverage the on-demand provisioning capabilities in your architecture, and work to refactor your application, in order to leverage the scalable infrastructure and take advantage of the cloud.

Characteristics of a truly scalable application:

Increasing resources results in a proportional increase in performance A scalable service is capable of handling heterogeneity

A scalable service is operationally efficient A scalable service is resilient

A scalable service should become more cost effective when it grows (Cost per unit reduces as the number of units increases)

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

NEW QUESTION 251

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 AM 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 255

You are playing around with setting up stacks using JSON templates in CloudFormation to try and understand them a little better. You have set up about 5 or 6 but now start to wonder if you are being charged for these stacks. What is AWS's billing policy regarding stack resources?

- A. You are not charged for the stack resources if they are not taking any traffic.
- B. You are charged for the stack resources for the time they were operating (even if you deleted the stack right away)
- C. You are charged for the stack resources for the time they were operating (but not if you deleted the stack within 60 minutes)
- D. You are charged for the stack resources for the time they were operating (but not if you deleted the stack within 30 minutes)

Answer: B

Explanation: A stack is a collection of AWS resources that you can manage as a single unit. In other words, you can create, update, or delete a collection of resources by creating, updating, or deleting stacks. All the resources in a stack are defined by the stack's AWS CloudFormation template. A stack, for instance, can include all the resources required to run a web application, such as a web server, a database, and networking rules. If you no longer require that web application, you can simply delete the stack, and all of its related resources are deleted.

You are charged for the stack resources for the time they were operating (even if you deleted the stack right away).

Reference: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacks.html>

NEW QUESTION 259

After deploying a new website for a client on AWS, he asks if you can set it up so that if it fails it can be automatically redirected to a backup website that he has stored on a dedicated server elsewhere. You are wondering whether Amazon Route 53 can do this. Which statement below is correct in regards to Amazon Route 53?

- A. Amazon Route 53 can't help detect an outage
- B. You need to use another service.
- C. Amazon Route 53 can help detect an outage of your website and redirect your end users to alternate locations.
- D. Amazon Route 53 can help detect an outage of your website but can't redirect your end users to alternate locations.
- E. Amazon Route 53 can't help detect an outage of your website, but can redirect your end users to alternate locations.

Answer: B

Explanation: With DNS Failover, Amazon Route 53 can help detect an outage of your website and redirect your end users to alternate locations where your application is operating properly.

Reference:

<http://aws.amazon.com/about-aws/whats-new/2013/02/11/announcing-dns-failover-for-route-53/>

NEW QUESTION 261

In Route 53, what does a Hosted Zone refer to?

- A. A hosted zone is a collection of geographical load balancing rules for Route 53.
- B. A hosted zone is a collection of resource record sets hosted by Route 53.
- C. A hosted zone is a selection of specific resource record sets hosted by CloudFront for distribution to Route 53.
- D. A hosted zone is the Edge Location that hosts the Route 53 records for a use

Answer: B

Explanation: A Hosted Zone refers to a selection of resource record sets hosted by Route 53.

Reference: <http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/AboutHostedZones.html>

NEW QUESTION 263

While creating a network in the VPC, which of the following is true of a NAT device?

- A. You have to administer the NAT Gateway Service provided by AWS.
- B. You can choose to use any of the three kinds of NAT devices offered by AWS for special purposes.
- C. You can use a NAT device to enable instances in a private subnet to connect to the Internet.
- D. You are recommended to use AWS NAT instances over NAT gateways, as the instances provide better availability and bandwidth.

Answer: C

Explanation: You can use a NAT device to enable instances in a private subnet to connect to the Internet (for example, for software updates) or other AWS services, but prevent the Internet from initiating connections with the instances. AWS offers two kinds of NAT devices: a NAT gateway or a NAT instance. We recommend NAT gateways, as they provide better availability and bandwidth over NAT instances. The NAT Gateway service is also a managed service that does not require your administration efforts. A NAT instance is launched from a NAT AM. You can choose to use a NAT instance for special purposes.

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

NEW QUESTION 268

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 stopped
- 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-interfaces-in-your-vpc>

NEW QUESTION 269

Your manager has asked you to set up a public subnet with instances that can send and receive internet traffic, and a private subnet that can't receive traffic directly from the internet, but can initiate traffic to the internet (and receive responses) through a NAT instance in the public subnet. Hence, the following 3 rules need to be allowed:

Inbound SSH traffic.

Web servers in the public subnet to read and write to MS SQL servers in the private subnet
Inbound RDP traffic from the Microsoft Terminal Services gateway in the public private subnet
What are the respective ports that need to be opened for this?

- A. Ports 22,1433,3389
- B. Ports 21,1433,3389
- C. Ports 25,1433,3389
- D. Ports 22,1343,3999

Answer: A

Explanation: A network access control list (ACL) is an optional layer of security that acts as a firewall for controlling traffic in and out of a subnet. You might set up network ACLs with rules similar to your security groups in order to add an additional layer of security to your VPC.

The following ports are recommended by AWS for a single subnet with instances that can receive and send Internet traffic and a private subnet that can't receive traffic directly from the Internet. However, it can initiate traffic to the Internet (and receive responses) through a NAT instance in the public subnet. Inbound SSH traffic. Port 22

Web servers in the public subnet to read and write to MS SQL servers in the private subnet. Port 1433 Inbound RDP traffic from the Microsoft Terminal Services gateway in the public private subnet. Port 3389 Reference:

http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Appendix_NACLs.html#VPC_Appendix_NACLs_Scenario_2

NEW QUESTION 274

You want to establish a dedicated network connection from your premises to AWS in order to save money by transferring data directly to AWS rather than through your internet service provider. You are sure there must be some other benefits beyond cost savings. Which of the following would not be considered a benefit if you were to establish such a connection?

- A. Elasticity
- B. Compatibility with all AWS services.
- C. Private connectivity to your Amazon VPC.
- D. Everything listed is a benefit

Answer: D

Explanation: AWS Direct Connect makes it easy to establish a dedicated network connection from your premises to AWS.

Using AWS Direct Connect, you can establish private connectivity between AWS and your datacenter, office, or colocation environment, which in many cases can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than internet-based connections.

You could expect the following benefits if you use AWS Direct Connect. Reduced bandwidth costs

Consistent network performance
Compatibility with all AWS services
Private connectivity to your Amazon VPC
Elasticity

Simplicity

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

NEW QUESTION 275

A user has launched an EC2 instance. The instance got terminated as soon as it was launched. Which of the below mentioned options is not a possible reason for this?

- A. The user account has reached the maximum volume limit
- B. The AMI is missing
- C. It is the required part
- D. The snapshot is corrupt
- E. The user account has reached the maximum EC2 instance limit

Answer: D

Explanation: When the user account has reached the maximum number of EC2 instances, it will not be allowed to launch an instance. AWS will throw an 'Instance Limit Exceeded' error. For all other reasons, such as "AMI is missing part", "Corrupt Snapshot" or "Volume limit has reached" it will launch an EC2 instance and then terminate it.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_InstanceStraightToTerminated.html

NEW QUESTION 279

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 manually join 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 280

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 284

Can you encrypt EBS volumes?

- A. Yes, you can enable encryption when you create a new EBS volume using the AWS Management Console, API, or CLI.
- B. No, you should use a third-party software to perform raw block-level encryption of an EBS volume.
- C. Yes, but you must use a third-party API for encrypting data before it's loaded on EBS.
- D. Yes, you can encrypt with the special "ebs_encrypt" command through Amazon API

Answer: A

Explanation: With Amazon EBS encryption, you can now create an encrypted EBS volume and attach it to a supported instance type. Data on the volume, disk I/O, and snapshots created from the volume are then all encrypted. The encryption occurs on the servers that host the EC2 instances, providing encryption of data as it moves between EC2 instances and EBS storage. EBS encryption is based on the industry standard AES-256 cryptographic algorithm.
To get started, simply enable encryption when you create a new EBS volume using the AWS Management Console, API, or CLI. Amazon EBS encryption is available for all the latest EC2 instances in all commercially available AWS regions.
Reference:
<https://aws.amazon.com/about-aws/whats-new/2014/05/21/Amazon-EBS-encryption-now-available/>

NEW QUESTION 288

A user has created an ELB with Auto Scaling. Which of the below mentioned offerings from ELB helps the user to stop sending new requests traffic from the load balancer to the EC2 instance when the instance is being deregistered while continuing in-flight requests?

- A. ELB sticky session

- B. ELB deregistration check
- C. ELB auto registration Off
- D. ELB connection draining

Answer: D

Explanation: The Elastic Load Balancer connection draining feature causes the load balancer to stop sending new requests to the back-end instances when the instances are deregistering or become unhealthy, while ensuring that in-flight requests continue to be served.

Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/config-conn-drain.html>

NEW QUESTION 290

A user is running a webserver on EC2. The user wants to receive the SMS when the EC2 instance utilization is above the threshold limit. Which AWS services should the user configure in this case?

- A. AWS CloudWatch + AWS SQS.
- B. AWS CloudWatch + AWS SNS.
- C. AWS CloudWatch + AWS SES.
- D. AWS EC2 + AWS Cloudwatc

Answer: B

Explanation: 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. In this case, the user can configure that Cloudwatch sends an alarm on when the threshold is crossed to SNS which will trigger an SMS.

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

NEW QUESTION 293

A user is making a scalable web application with compartmentalization. The user wants the log module to be able to be accessed by all the application functionalities in an asynchronous way. Each module of the application sends data to the log module, and based on the resource availability it will process the logs. Which AWS service helps this functionality?

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

Answer: A

Explanation: Amazon Simple Queue Service (SQS) is a highly reliable distributed messaging system for storing messages as they travel between computers. By using Amazon SQS, developers can simply move data between distributed application components. It is used to achieve compartmentalization or loose coupling. In this case all the modules will send a message to the logger queue and the data will be processed by queue as per the resource availability.

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

NEW QUESTION 295

Your manager has come to you saying that he is very confused about the bills he is receiving from AWS as he is getting different bills for every user and needs you to look into making it more understandable. Which of the following would be the best solution to meet his request?

- A. AWS Billing Aggregation
- B. Consolidated Billing
- C. Deferred Billing
- D. Aggregated Billing

Answer: B

Explanation: Consolidated Billing enables you to consolidate payment for multiple AWS accounts within your company by designating a single paying account. Consolidated Billing enables you to see a combined view of AWS costs incurred by all accounts, as well as obtain a detailed cost report for each of the individual AWS accounts associated with your "Paying Account". Consolidated Billing is offered at no additional charge. Reference: <https://aws.amazon.com/billing/faqs/>

NEW QUESTION 297

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 Mobile 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 299

You have been asked to set up a database in AWS that will require frequent and granular updates. You know that you will require a reasonable amount of storage space but are not sure of the best option. What is the recommended storage option when you run a database on an instance with the above criteria?

- A. Amazon S3
- B. Amazon EBS
- C. AWS Storage Gateway
- D. Amazon Glacier

Answer: B

Explanation: Amazon EBS provides durable, block-level storage volumes that you can attach to a running Amazon EC2 instance. You can use Amazon EBS as a primary storage device for data that requires frequent and granular updates. For example, Amazon EBS is the recommended storage option when you run a database on an instance.

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

NEW QUESTION 300

You have been asked to set up monitoring of your network and you have decided that Cloudwatch would be the best service to use. Amazon CloudWatch monitors your Amazon Web Services (AWS) resources and the applications you run on AWS in real-time. You can use CloudWatch to collect and track metrics, which are the variables you want to measure for your resources and applications. Which of the following items listed can AWS Cloudwatch monitor?

- A. Log files your applications generate.
- B. All of the items listed on this page.
- C. System-wide visibility into resource utilization, application performance, and operational health.
- D. Custom metrics generated by your applications and services .

Answer: B

Explanation: Amazon CloudWatch can monitor AWS resources such as Amazon EC2 instances, Amazon DynamoDB tables, and Amazon RDS DB instances, as well as custom metrics generated by your applications and services, and any log files your applications generate. You can use Amazon CloudWatch to gain system-wide visibility into resource utilization, application performance, and operational health. You can use these insights to react and keep your application running smoothly.

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

NEW QUESTION 301

A user has hosted an application on EC2 instances. The EC2 instances are configured with ELB and Auto Scaling. The application server session time out is 2 hours. The user wants to configure connection draining to ensure that all in-flight requests are supported by ELB even though the instance is being deregistered. What time out period should the user specify for connection draining?

- A. 1 hour
- B. 30 minutes
- C. 5 minutes
- D. 2 hours

Answer: A

Explanation: The Elastic Load Balancer connection draining feature causes the load balancer to stop sending new requests to the back-end instances when the instances are deregistering or become unhealthy, while ensuring that in-flight requests continue to be served. The user can specify a maximum time of 3600 seconds (1 hour) for the load balancer to keep the connections alive before reporting the instance as deregistered. If the user does not specify the maximum timeout period, by default, the load balancer will close the connections to the deregistering instance after 300 seconds.

Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/config-conn-drain.html>

NEW QUESTION 305

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 308

Identify a true statement about the On-Demand instances purchasing option provided by Amazon EC2.

- A. Pay for the instances that you use by the hour, with no long-term commitments or up-front payments.
- B. Make a low, one-time, up-front payment for an instance, reserve it for a one- or three-year term, and pay a significantly lower hourly rate for these instances.
- C. Pay for the instances that you use by the hour, with long-term commitments or up-front payments.

D. Make a high, one-time, all-front payment for an instance, reserve it for a one- or three-year term, and pay a significantly higher hourly rate for these instance

Answer: A

Explanation: On-Demand instances allow you to pay for the instances that you use by the hour, with no long-term commitments or up-front payments.
Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/reserved-instances-offerings.html>

NEW QUESTION 309

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 314

Your customer wishes to deploy an enterprise application to AWS which will consist of several web servers, several application servers and a small (50GB) Oracle database information is stored, both in the database and the file systems of the various servers. The backup system must support database recovery whole server and whole disk restores, and individual file restores with a recovery time of no more than two hours. They have chosen to use RDS Oracle as the database Which backup architecture will meet these requirements?

- A. Backup RDS using automated daily DB backups Backup the EC2 instances using AMIs and supplement with file-level backup to S3 using traditional enterprise backup software to provide file level restore
- B. Backup RDS using a Multi-AZ Deployment Backup the EC2 instances using AMIs, and supplement by copying file system data to S3 to provide file level restore.
- C. Backup RDS using automated daily DB backups Backup the EC2 instances using EBS snapshots and supplement with file-level backups to Amazon Glacier using traditional enterprise backup software to provide file level restore
- D. Backup RDS database to S3 using Oracle RMAN Backup the EC2 instances using AMIs, and supplement with EBS snapshots for individual volume restore.

Answer: A

Explanation: Point-In-Time Recovery

In addition to the daily automated backup, Amazon RDS archives database change logs. This enables you to recover your database to any point in time during the backup retention period, up to the last five minutes of database usage.

Amazon RDS stores multiple copies of your data, but for Single-AZ DB instances these copies are stored in a single availability zone. If for any reason a Single-AZ DB instance becomes unusable, you can use point-in-time recovery to launch a new DB instance with the latest restorable data. For more information on working with point-in-time recovery, go to Restoring a DB Instance to a Specified Time.

Note

Multi-AZ deployments store copies of your data in different Availability Zones for greater levels of data durability. For more information on Multi-AZ deployments, see High Availability (Multi-AZ).

NEW QUESTION 318

Your company has HQ in Tokyo and branch offices all over the world and is using a logistics software with a multi-regional deployment on AWS in Japan, Europe and USA, The logistic software has a 3-tier architecture and currently uses MySQL 5.6 for data persistence. Each region has deployed its own database In the HQ region you run an hourly batch process reading data from every region to compute cross regional reports that are sent by email to all offices this batch process must be completed as fast as possible to quickly optimize logistics how do you build the database architecture in order to meet the requirements'?

- A. For each regional deployment, use RDS MySQL with a master in the region and a read replica in the HQ region
- B. For each regional deployment, use MySQL on EC2 with a master in the region and send hourly EBS snapshots to the HQ region
- C. For each regional deployment, use RDS MySQL with a master in the region and send hourly RDS snapshots to the HQ region
- D. For each regional deployment, use MySQL on EC2 with a master in the region and use S3 to copy data files hourly to the HQ region
- E. Use Direct Connect to connect all regional MySQL deployments to the HQ region and reduce network latency for the batch process

Answer: A

NEW QUESTION 320

A customer has a 10 GB AWS Direct Connect connection to an AWS region where they have a web application hosted on Amazon Elastic Computer Cloud (EC2). The application has dependencies on an on-premises mainframe database that uses a BASE (Basic Available. Sort stale Eventual consistency) rather than an ACID (Atomicity. Consistency isolation. Durability) consistency model.

The application is exhibiting undesirable behavior because the database is not able to handle the volume of writes. How can you reduce the load on your on-premises database resources in the most cost-effective way?

- A. Use an Amazon Elastic Map Reduce (EMR) S3DistCp as a synchronization mechanism between the on-premises database and a Hadoop cluster on AWS.
- B. Modify the application to write to an Amazon SQS queue and develop a worker process to flush the queue to the on-premises database.
- C. Modify the application to use DynamoDB to feed an EMR cluster which uses a map function to write to the on-premises database.
- D. Provision an RDS read-replica database on AWS to handle the writes and synchronize the two databases using Data Pipeline.

Answer: A

Explanation: Reference: <https://aws.amazon.com/blogs/aws/category/amazon-elastic-map-reduce/>

NEW QUESTION 323

Your company plans to host a large donation website on Amazon Web Services (AWS). You anticipate a large and undetermined amount of traffic that will create many database writes. To be certain that you do not drop any writes to a database hosted on AWS. Which service should you use?

- A. Amazon RDS with provisioned IOPS up to the anticipated peak write throughput.
- B. Amazon Simple Queue Service (SQS) for capturing the writes and draining the queue to write to the database.
- C. Amazon ElastiCache to store the writes until the writes are committed to the database.
- D. Amazon DynamoDB with provisioned write throughput up to the anticipated peak write throughput

Answer: B

Explanation: Amazon Simple Queue Service (Amazon SQS) offers a reliable, highly scalable hosted queue for storing messages as they travel between computers. By using Amazon SQS, developers can simply move data between distributed application components performing different tasks, without losing messages or requiring each component to be always available. Amazon SQS makes it easy to build a distributed, decoupled application, working in close conjunction with the Amazon Elastic Compute Cloud (Amazon EC2) and the other AWS infrastructure web services.

What can I do with Amazon SQS?

Amazon SQS is a web service that gives you access to a message queue that can be used to store messages while waiting for a computer to process them. This allows you to quickly build message queuing applications that can be run on any computer on the internet. Since Amazon SQS is highly scalable and you only pay for what you use, you can start small and grow your application as you wish, with no compromise on performance or reliability. This lets you focus on building sophisticated message-based applications, without worrying about how the messages are stored and managed.

You can use Amazon SQS with software applications in various ways. For example, you can: Integrate Amazon SQS with other AWS infrastructure web services to make applications more reliable and flexible.

Use Amazon SQS to create a queue of work where each message is a task that needs to be completed by a process. One or many computers can read tasks from the queue and perform them. Build a microservices architecture, using queues to connect your microservices.

Keep notifications of significant events in a business process in an Amazon SQS queue. Each event can have a corresponding message in a queue, and applications that need to be aware of the event can read and process the messages.

NEW QUESTION 327

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 volume storage 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 331

You have recently joined a startup company building sensors to measure street noise and air quality in urban areas. The company has been running a pilot deployment of around 100 sensors for 3 months each sensor uploads 1KB of sensor data every minute to a backend hosted on AWS.

During the pilot, you measured a peak of 10 IOPS on the database, and you stored an average of 3GB of sensor data per month in the database.

The current deployment consists of a load-balanced auto scaled ingestion layer using EC2 instances and a PostgreSQL RDS database with 500GB standard storage.

The pilot is considered a success and your CEO has managed to get the attention of some potential investors. The business plan requires a deployment of at least 1,000 sensors which needs to be supported by the backend. You also need to store sensor data for at least two years to be able to compare year over year improvements.

To secure funding, you have to make sure that the platform meets these requirements and leaves room for further scaling. Which setup will meet the requirements?

- A. Add an SQS queue to the ingestion layer to buffer writes to the RDS instance
- B. Ingest data into a DynamoDB table and move old data to a Redshift cluster
- C. Replace the RDS instance with a 6 node Redshift cluster with 96TB of storage
- D. Keep the current architecture but upgrade RDS storage to 3TB and IOPS provisioned IOPS

Answer: C

NEW QUESTION 334

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 339

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 344

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 349

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 Dynamo DB 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 item into an SQS queue in the second region; use an auto-scaling group behind the SQS queue to replay the write in the second region.

Answer: A

NEW QUESTION 350

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 priority and verify network traffic is leveraging the DirectConnect connection.
- E. And verify network traffic is leveraging the DirectConnect connection.

Answer: D

NEW QUESTION 354

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 Cloud Front 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 and health checks and DNS failover.
- 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 358

You are tasked with moving a legacy application from a virtual machine running inside your datacenter to an Amazon VPC Unfortunately this app requires access to a number of on-premises services and no one who configured the app still works for your company. Even worse there's no documentation for it. What will allow the application running inside the VPC to reach back and access its internal dependencies without being reconfigured? (Choose 3 answers)

- A. An AWS Direct Connect link between the VPC and the network housing the internal services.
- B. An Internet Gateway to allow a VPN connection.
- C. An Elastic IP address on the VPC instance
- D. An IP address space that does not conflict with the one on-premises
- E. Entries in Amazon Route 53 that allow the Instance to resolve its dependencies' IP addresses
- F. A VM Import of the current virtual machine

Answer: ADF

Explanation: AWS Direct Connect

AWS Direct Connect makes it easy to establish a dedicated network connection from your premises to AWS. Using AWS Direct Connect, you can establish private connectivity between AWS your datacenter, office, or colocation environment, which in many cases can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than Internet based connections.

AWS Direct Connect lets you establish a dedicated network connection between your network and one of the AWS Direct Connect locations. Using industry standard 802.1q VLANs, this dedicated connection can be partitioned into multiple virtual interfaces. This allows you to use the same connection to access public resources such as objects stored in Amazon S3 using public IP address space, and private resources such as Amazon EC2 instances running within an Amazon Virtual Private Cloud (VPC) using private IP space, while maintaining network separation between the public and private environments. Virtual interfaces can be reconfigured at any time to meet your changing needs.

What is AWS Direct Connect?

AWS Direct Connect links your internal network to an AWS Direct Connect location over a standard 1 gigabit or 10 gigabit Ethernet fiber-optic cable. One end of the cable is connected to your router, the other to an AWS Direct Connect router. With this connection in place, you can create virtual interfaces directly to the AWS cloud (for example, to Amazon Elastic Compute Cloud (Amazon EC2) and Amazon Simple Storage Service (Amazon S3)) and to Amazon Virtual Private Cloud (Amazon VPC), bypassing Internet service providers in your network path. An AWS Direct Connect location provides access to Amazon Web Services in the region it is associated with, as well as access to other US regions. For example, you can provision a single connection to any AWS Direct Connect location in the US and use it to access public AWS services in all US Regions and AWS GovCloud (US).

The following diagram shows how AWS Direct Connect interfaces with your network.

Requirements

To use AWS Direct Connect, your network must meet one of the following conditions:

Your network is colocated with an existing AWS Direct Connect location. For more information on available AWS Direct Connect locations, go to

http://aws.amazon.com/directconnect/.

You are working with an AWS Direct Connect partner who is a member of the AWS Partner Network (APN). For a list of AWS Direct Connect partners who can help you connect, go to <http://aws.amazon.com/directconnect>

You are working with an independent service provider to connect to AWS Direct Connect. In addition, your network must meet the following conditions:

Connections to AWS Direct Connect require single mode fiber, 1000BASE-LX (1310nm) for 1 gigabit Ethernet, or 10GBASE-LR (1310nm) for 10 gigabit Ethernet. Auto Negotiation for the port must be disabled. You must support 802.1Q VLANs across these connections.

Your network must support Border Gateway Protocol (BGP) and BGP MD5 authentication. Optionally, you may configure Bidirectional Forwarding Detection (BFD).

To connect to Amazon Virtual Private Cloud (Amazon VPC), you must first do the following: Provide a private Autonomous System Number (ASN). Amazon allocates a private IP address in the 169.x.x.x range to you.

Create a virtual private gateway and attach it to your VPC. For more information about creating a virtual private gateway, see Adding a Hardware Virtual Private Gateway to Your VPC in the Amazon VPC User Guide.

To connect to public AWS products such as Amazon EC2 and Amazon S3, you need to provide the following:

A public ASN that you own (preferred) or a private ASN.

Public IP addresses (/31) (that is, one for each end of the BGP session) for each BGP session. If you do not have public IP addresses to assign to this connection, log on to AWS and then open a ticket with AWS Support.

The public routes that you will advertise over BGP.

NEW QUESTION 359

A newspaper organization has a on-premises application which allows the public to search its back catalogue and retrieve individual newspaper pages via a website written in Java. They have scanned the old newspapers into JPEGs (approx 17TB) and used Optical Character Recognition (OCR) to populate a commercial search product. The hosting platform and software are now end of life and the organization wants to migrate its archive to AWS and produce a cost efficient architecture and still be designed for availability and durability. Which is the most appropriate?

- A. Use S3 with reduced redundancy to store and serve the scanned files, install the commercial search application on EC2 Instances and configure with auto-scaling and an Elastic Load Balancer.
- B. Model the environment using CloudFormation use an EC2 instance running Apache webserver and an open source search application, stripe multiple standard EB5 volumes together to store the JPEGs and search index.
- C. Use S3 with standard redundancy to store and serve the scanned files, use CloudSearch for queryprocessing, and use Elastic Beanstalk to host the website across multiple availability zones.
- D. Use a single-AZ RD5 My5QL instance to store the search index and the JPEG images use an EC2 instance to serve the website and translate user queries into SQL.
- E. Use a CloudFront download distribution to serve the JPEGs to the end users and Install the current commercial search product, along with a Java Container on the website on EC2 instances and use Route53 with DNS round-robin.

Answer: C

Explanation: There is no such thing as "Most appropriate" without knowing all your goals. I find your scenarios very fuzzy, since you can obviously mix-n-match between them. I think you should decide by layers instead: Load Balancer Layer: ELB or just DNS, or roll-your-own. (Using DNS+ELB is slightly cheaper, but less reliable than ELB.)

Storage Layer for 17TB of Images: This is the perfect use case for S3. Off-load all the web requests directly to the relevant JPEGs in S3. Your EC2 boxes just generate links to them.

If your app already serves its own images (not links to images), you might start with EFS. But more than likely, you can just setup a web server to re-write or re-direct all JPEG links to S3 pretty easily.

If you use S3, don't serve directly from the bucket- Serve via a CNAME in domain you control. That way, you can switch in CloudFront easily.

EBS will be way more expensive, and you'll need 2x the drives if you need 2 boxes. Yuck. Consider a smaller storage format. For example, JPEG200 or WebP or other tools might make for smaller images. There is also the DejaVu format from a while back.

Cache Layer: Adding Cloud Front in front of S3 will help people on the other side of the world-- well, possibly. Typical archives follow a power law. The long tail of requests means that most JPEGs won't be requested enough to be in the cache. So you are only speeding up the most popular objects. You can always wait, and switch in CF later after you know your costs better. (In some cases, it can actually lower costs.)

You can also put CloudFront in front of your app, since your archive search results should be fairly static. This will also allow you to run with a smaller instance type, since CF will handle much of the load if you do it right.

Database Layer: A few options:

Use whatever your current server does for now, and replace with something else down the road. Don't under-estimate this approach, sometimes it's better to start now and optimize later.

Use RDS to run MySQL/ Postgres

I'm not as familiar with ElasticSearch or Cloudsearch, but obviously Cloudsearch will be less maintenance+setup.

App Layer:

When creating the app layer from scratch, consider Cloud Formation and/or OpsWorks. It's extra stuff to learn, but helps down the road.

Java+ Tomcat is right up the alley of ElasticBeanstalk. (Basically EC2 + Autoscale + ELB).

Preventing Abuse: When you put something in a public S3 bucket, people will hot-link it from their web pages. If you want to prevent that, your app on the EC2 box can generate signed links to S3 that expire in a few hours. Now everyone will be forced to go thru the app, and the app can apply rate limiting, etc. Saving money:

If you don't mind having downtime:

run everything in one AZ (both DBs and EC2s). You can always add servers and AZs down the road, as long as it's architected to be stateless. In fact, you should use multiple regions if you want it to be really robust.

use Reduced Redundancy in S3 to save a few hundred bucks per month (Someone will have to "go fix it" every time it breaks, including having an off-line copy to repair S3.)

Buy Reserved Instances on your EC2 boxes to make them cheaper. (Start with the RI market and buy a partially used one to get started.) It's just a coupon saying "if you run this type of box in this AZ, you will save on the per-hour costs." You can get 1/2 to 1/3 off easily.

Rewrite the application to use less memory and CPU -that way you can run on fewer/ smaller boxes. (May or may not be worth the investment.)

If your app will be used very infrequently, you will save a lot of money by using Lambda. I'd be worried that it would be quite slow if you tried to run a Java application on it though ..

We're missing some information like load, latency expectations from search, indexing speed, size of the search index, etc. But with what you've given us, I would go with S3 as the storage for the files (S3 rocks. It is really, really awesome). If you're stuck with the commercial search application, then on EC2 instances with autoscaling and an ELB. If you are allowed an alternative search engine, Elasticsearch is probably your best bet. I'd run it on EC2 instead of the AWS Elasticsearch service, as IMHO it's not ready yet. Don't autoscale Elasticsearch automatically though, it'll cause all sorts of issues. I have zero experience with CloudSearch so I can't comment on that. Regardless of which option, I'd use Cloud Formation for all of it.

NEW QUESTION 363

Your company has an on-premises multi-tier PHP web application, which recently experienced downtime due to a large burst in web traffic due to a company announcement. Over the coming days, you are expecting similar announcements to drive similar unpredictable bursts, and are looking to find ways to quickly improve your infrastructure's ability to handle unexpected increases in traffic.

The application currently consists of 2 tiers: a web tier which consists of a load balancer and several Linux Apache web servers as well as a database tier which hosts a Linux server hosting a MySQL database. Which scenario below will provide full site functionality, while helping to improve the ability of your application in the short timeframe required?

- A. Failover environment: Create an S3 bucket and configure it for website hosting.
- B. Migrate your DNS to Route53 using zone file import, and leverage Route53 DNS failover to failover to the S3 hosted website.
- C. Hybrid environment: Create an AMI, which can be used to launch web servers in EC2. Create an Auto Scaling group, which uses the AMI to scale the web tier based on incoming traffic.
- D. Leverage Elastic Load Balancing to balance traffic between on-premises web servers and those hosted in AWS.
- E. Offload traffic from on-premises environment: Setup a CloudFront distribution, and configure CloudFront to cache objects from a custom origin.
- F. Choose to customize your object cache behavior, and select a TTL that objects should exist in cache.
- G. Migrate to AWS: Use VM Import/Export to quickly convert an on-premises web server to an AMI.
- H. Create an Auto Scaling group, which uses the imported AMI to scale the web tier based on incoming traffic.
- I. Create an RDS read replica and setup replication between the RDS instance and on-premises MySQL server to migrate the database.

Answer: C

NEW QUESTION 368

Your company produces customer-commissioned one-of-a-kind skiing helmets combining high fashion with custom technical enhancements. Customers can show off their individuality on the ski slopes and have access to head-up-displays, GPS rear-view cams, and any other technical innovation they wish to embed in the helmet.

The current manufacturing process is data-rich and complex, including assessments to ensure that the custom electronics and materials used to assemble the helmets are to the highest standards. Assessments are a mixture of human and automated assessments; you need to add a new set of assessment to model the failure modes of the custom electronics using GPUs with CUDA, across a cluster of servers with low latency networking.

What architecture would allow you to automate the existing process using a hybrid approach and ensure that the architecture can support the evolution of processes over time?

- A. Use AWS Data Pipeline to manage movement of data & meta-data and assessments. Use an auto-scaling group of G2 instances in a placement group.
- B. Use Amazon Simple Workflow (SWF) to manage assessments, movement of data & meta-data. Use an auto-scaling group of G2 instances in a placement group.
- C. Use Amazon Simple Workflow (SWF) to manage assessments, movement of data & meta-data. Use an auto-scaling group of C3 instances with SR-IOV (Single Root I/O Virtualization).
- D. Use AWS Data Pipeline to manage movement of data & meta-data and assessments. Use an auto-scaling group of C3 instances with SR-IOV (Single Root I/O virtualization).

Answer: B

NEW QUESTION 372

You require the ability to analyze a large amount of data, which is stored on Amazon S3 using Amazon Elastic MapReduce. You are using the cc2.8xlarge instance type, whose CPUs are mostly idle during processing. Which of the below would be the most cost-efficient way to reduce the runtime of the job?

- A. Create more smaller files on Amazon S3.
- B. Add additional cc2.8xlarge instances by introducing a task group.
- C. Use smaller instances that have higher aggregate I/O performance.
- D. Create fewer, larger files on Amazon S3.

Answer: C

NEW QUESTION 374

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 job.
- 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 job.
- 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 378

Your website is serving on-demand training videos to your workforce. Videos are uploaded monthly in high resolution MP4 format. Your workforce is distributed globally often on the move and using company-provided tablets that require the HTTP Live Streaming (HLS) protocol to watch a video. Your company has no video transcoding expertise and it required you may need to pay for a consultant.

How do you implement the most cost-efficient architecture without compromising high availability and quality of video delivery?

- A. A video transcoding pipeline running on EC2 using SQS to distribute tasks and Auto Scaling to adjust the number of nodes depending on the length of the queue
- B. EBS volumes to host videos and EBS snapshots to incrementally backup original files after a few days
- C. CloudFront to serve HLS transcoded videos from EC2.
- D. Elastic Transcoder to transcode original high-resolution MP4 videos to H.264
- E. EBS volumes to host videos and EBS snapshots to incrementally backup original files after a few days
- F. CloudFront to serve HLS transcoded videos from EC2.
- G. Elastic Transcoder to transcode original high-resolution MP4 videos to H.264
- H. S3 to host videos with Lifecycle Management to archive original files to Glacier after a few days
- I. CloudFront to serve HLS transcoded videos from S3.
- J. A video transcoding pipeline running on EC2 using SQS to distribute tasks and Auto Scaling to adjust the number of nodes depending on the length of the queue
- K. S3 to host videos with Lifecycle Management to archive all files to Glacier after a few days
- L. CloudFront to serve HLS transcoded videos from Glacier.

Answer: C

NEW QUESTION 381

An AWS customer is deploying an application that is composed of an AutoScaling group of EC2 Instances.

The customer's security policy requires that every outbound connection from these instances to any other service within the customer's Virtual Private Cloud must be authenticated using a unique X.509 certificate that contains the specific instance-id.

In addition, X.509 certificates must be designed by the customer's Key Management Service in order to be trusted for authentication.

Which of the following configurations will support these requirements?

- A. Configure an IAM Role that grants access to an Amazon S3 object containing a signed certificate and configure the Auto Scaling group to launch instances with this role. Have the instances bootstrap get the certificate from Amazon S3 upon first boot.
- B. Embed a certificate into the Amazon Machine Image that is used by the Auto Scaling group. Have the launched instances generate a certificate signature request with the instance's assigned instance-id to the Key Management Service for signature.
- C. Configure the Auto Scaling group to send an SNS notification of the launch of a new instance to the trusted Key Management Service.
- D. Have the Key Management Service generate a signed certificate and send it directly to the newly launched instance.
- E. Configure the launched instances to generate a new certificate upon first boot. Have the Key Management Service poll the AutoScaling group for associated instances and send new instances a certificate signature (that contains the specific instance-id).

Answer: A

NEW QUESTION 384

You are designing an S3 solution that requires HTTPS clients to be authenticated by the Web server using client certificate authentication. The solution must be resilient.

Which of the following options would you consider for configuring the web server infrastructure? (Choose 2 answers)

- A. Configure ELB with TCP listeners on TCP/443. And place the Web servers behind it.
- B. Configure your Web servers with HTTPS. Place the Web servers in a Route53 Record Set and configure health checks against all Web servers.
- C. Configure ELB with HTTPS listeners, and place the Web servers behind it.
- D. Configure your web servers as the origins for a CloudFront distribution.
- E. Use custom SSL certificates on your CloudFront distribution.

Answer: AB

NEW QUESTION 388

Your Fortune 500 company has undertaken a TCO analysis evaluating the use of Amazon S3 versus acquiring more hardware. The outcome was that all employees would be granted access to use Amazon S3 for storage of their personal documents.

Which of the following will you need to consider so you can set up a solution that incorporates single sign-on from your corporate AD or LDAP directory and restricts access for each user to a designated user folder in a bucket? (Choose 3 Answers)

- A. Setting up a federation proxy or identity provider
- B. Using AWS Security Token Service to generate temporary tokens
- C. Tagging each folder in the bucket
- D. Configuring IAM role
- E. Setting up a matching IAM user for every user in your corporate directory that needs access to a folder in the bucket

Answer: ABD

NEW QUESTION 393

You require the ability to analyze a customer's clickstream data on a website so they can do behavioral analysis. Your customer needs to know what sequence of pages and ads their customer clicked on. This data will be used in real time to modify the page layouts as customers click through the site to increase stickiness.

and advertising click-through. Which option meets the requirements for captioning and analyzing this data?

- A. Log clicks in weblogs by URL store to Amazon S3, and then analyze with Elastic MapReduce
- B. Push web clicks by session to Amazon Kinesis and analyze behavior using Kinesis workers
- C. Write click events directly to Amazon Redshift and then analyze with SQL
- D. Publish web clicks by session to an Amazon SQS queue periodically drain these events to Amazon RDS and analyze with SQL

Answer: B

Explanation: Reference: <http://www.slideshare.net/AmazonWebServices/aws-webcast-introduction-to-amazon-kinesis>

NEW QUESTION 397

An AWS customer runs a public blogging website. The site users upload two million blog entries a month. The average blog entry size is 200 KB. The access rate to blog entries drops to negligible 6 months after publication and users rarely access a blog entry 1 year after publication. Additionally, blog entries have a high update rate during the first 3 months following publication, this drops to no updates after 6 months. The customer wants to use CloudFront to improve his user's load times.

Which of the following recommendations would you make to the customer?

- A. Duplicate entries into two different buckets and create two separate CloudFront distributions where S3 access is restricted only to Cloud Front identity
- B. Create a CloudFront distribution with "US" Europe price class for US/ Europe users and a different CloudFront distribution with All Edge Locations' for the remaining users.
- C. Create a CloudFront distribution with S3 access restricted only to the CloudFront identity and partition the blog entry's location in S3 according to the month it was uploaded to be used with CloudFront behaviors.
- D. Create a CloudFront distribution with Restrict Viewer Access Forward Query string set to true and minimum TTL of 0.

Answer: C

NEW QUESTION 398

A company is running a batch analysis every hour on their main transactional DB. running on an RDS MySQL instance to populate their central Data Warehouse running on Redshift. During the execution of the batch their transactional applications are very slow. When the batch completes they need to update the top management dashboard with the new data. The dashboard is produced by another system running on-premises that is currently started when a manually-sent email notifies that an update is required. The on-premises system cannot be modified because it is managed by another team.

How would you optimize this scenario to solve performance issues and automate the process as much as possible?

- A. Replace RDS with Redshift for the batch analysis and SNS to notify the on-premises system to update the dashboard
- B. Replace RDS with Redshift for the batch analysis and SQS to send a message to the on-premises system to update the dashboard
- C. Create an RDS Read Replica for the batch analysis and SNS to notify the on-premises system to update the dashboard
- D. Create an RDS Read Replica for the batch analysis and SQS to send a message to the on-premises system to update the dashboard.

Answer: A

NEW QUESTION 403

You are implementing a URL whitelisting system for a company that wants to restrict outbound HTTP'S connections to specific domains from their EC2-hosted applications. You deploy a single EC2 instance running proxy software and configure it to accept traffic from all subnets and EC2 instances in the VPC. You configure the proxy to only pass through traffic to domains that you define in its whitelist configuration. You have a nightly maintenance window of 10 minutes where all instances fetch new software updates. Each update is about 200MB in size and there are 500 instances in the VPC that routinely fetch updates. After a few days you notice that some machines are failing to successfully download some, but not all of their updates within the maintenance window. The download URLs used for these updates are correctly listed in the proxy's whitelist configuration and you are able to access them manually using a web browser on the instances. What might be happening? (Choose 2 answers)

- A. You are running the proxy on an undersized EC2 instance type so network throughput is not sufficient for all instances to download their updates in time.
- B. You are running the proxy on a sufficiently-sized EC2 instance in a private subnet and its network throughput is being throttled by a NAT running on an undersized EC2 instance.
- C. The route table for the subnets containing the affected EC2 instances is not configured to direct network traffic for the software update locations to the proxy.
- D. You have not allocated enough storage to the EC2 instance running the proxy so the network buffer is filling up, causing some requests to fail.
- E. You are running the proxy in a public subnet but have not allocated enough EIPs to support the needed network throughput through the Internet Gateway (IGW).

Answer: AB

NEW QUESTION 408

To serve Web traffic for a popular product your chief financial officer and IT director have purchased 10 ml large heavy utilization Reserved Instances (RIs) evenly spread across two availability zones:

Route 53 is used to deliver the traffic to an Elastic Load Balancer (ELB). After several months, the product grows even more popular and you need additional capacity. As a result, your company purchases two C3.2xlarge medium utilization RIs. You register the two c3 2xlarge instances with your ELB and quickly find that the ml large instances are at 100% of capacity and the c3 2xlarge instances have significant capacity that's unused. Which option is the most cost effective and uses EC2 capacity most effectively?

- A. Use a separate ELB for each instance type and distribute load to ELBs with Route 53 weighted round robin
- B. Configure Autoscaling group and Launch Configuration with ELB to add up to 10 more on-demand ml large instances when triggered by Cloudwatch shut off c3 2xlarge instances
- C. Route traffic to EC2 ml large and c3 2xlarge instances directly using Route 53 latency based routing and health checks shut off ELB
- D. Configure ELB with two c3 2xlarge Instances and use on-demand Autoscaling group for up to two additional c3.2xlarge instances. Shut on ml .large instances.

Answer: D

NEW QUESTION 410

A large real-estate brokerage is exploring the option of adding a cost-effective location based alert to their existing mobile application. The application backend infrastructure currently runs on AWS. Users who opt in to this service will receive alerts on their mobile device regarding real-estate offers in proximity to their location. For the alerts to be relevant, delivery time needs to be in the low minute count. The existing mobile app has 5 million users across the US. Which one of the following architectural suggestions would you make to the customer?

- A. The mobile application will submit its location to a web service endpoint utilizing Elastic Load Balancing and EC2 instances. DynamoDB will be used to store and retrieve relevant offers. EC2 instances will communicate with mobile carriers/device providers to push alerts back to mobile application.
- B. Use AWS DirectConnect or VPN to establish connectivity with mobile carriers. EC2 instances will receive the mobile applications' location through carrier connection. ROS will be used to store and relevant offers. EC2 instances will communicate with mobile carriers to push alerts back to the mobile application.
- C. The mobile application will send device location using SNS.
- D. EC2 instances will retrieve the relevant offers from DynamoDB. AWS Mobile Push will be used to send offers to the mobile application.
- E. The mobile application will send device location using AWS Mobile Push. EC2 instances will retrieve the relevant offers from DynamoDB. EC2 instances will communicate with mobile carriers/device providers to push alerts back to the mobile application.

Answer: A

NEW QUESTION 411

Your team has a tomcat-based Java application you need to deploy into development, test and production environments. After some research, you opt to use Elastic Beanstalk due to its tight integration with your developer tools and RDS due to its ease of management. Your QA team lead points out that you need to roll a sanitized set of production data into your environment on a nightly basis. Similarly, other software teams in your org want access to that same restored data via their EC2 instances in your VPC. The optimal setup for persistence and security that meets the above requirements would be the following.

- A. Create your RDS instance as part of your Elastic Beanstalk definition and alter its security group to allow access to it from hosts in your application subnets.
- B. Create your RDS instance separately and add its IP address to your application's DB connection strings in your code. Alter its security group to allow access to it from hosts within your VPC's IP address block.
- C. Create your RDS instance separately and pass its DNS name to your app's DB connection string as an environment variable.
- D. Create a security group for client machines and add it as a valid source for DB traffic to the security group of the RDS instance itself.
- E. Create your RDS instance separately and pass its DNS name to your app's DB connection string as an environment variable. Alter its security group to allow access to it from hosts in your application subnets.

Answer: A

NEW QUESTION 414

Your customer is willing to consolidate their log streams (access logs, application logs, security logs, etc.) in one single system. Once consolidated, the customer wants to analyze these logs in real time based on heuristics. From time to time, the customer needs to validate heuristics, which requires going back to data samples extracted from the last 12 hours?

What is the best approach to meet your customer's requirements?

- A. Send all the log events to Amazon SQS.
- B. Setup an Auto Scaling group of EC2 instances to consume the logs and apply the heuristics.
- C. Send all the log events to Amazon Kinesis. Develop a client process to apply heuristics on the logs.
- D. Configure Amazon Cloud Trail to receive custom logs, use EMR to apply heuristics on the logs.
- E. Setup an Auto Scaling group of EC2 syslogd servers, store the logs on S3, use EMR to apply heuristics on the logs.

Answer: B

Explanation: The throughput of an Amazon Kinesis stream is designed to scale without limits via increasing the number of shards within a stream. However, there are certain limits you should keep in mind while using Amazon Kinesis Streams:

By default, records of a stream are accessible for up to 24 hours from the time they are added to the stream. You can raise this limit to up to 7 days by enabling extended data retention.

The maximum size of a data blob (the data payload before Base64-encoding) within one record is 1 megabyte (MB).

Each shard can support up to 1000 PUT records per second.

For more information about other API level limits, see Amazon Kinesis Streams Limits.

NEW QUESTION 419

A web company is looking to implement an intrusion detection and prevention system into their deployed VPC. This platform should have the ability to scale to thousands of instances running inside of the VPC. How should they architect their solution to achieve these goals?

- A. Configure an instance with monitoring software and the elastic network interface (ENI) set to promiscuous mode. Packet sniffing to see all traffic across the VPC.
- B. Create a second VPC and route all traffic from the primary application VPC through the second VPC where the scalable virtualized IDS/IPS platform resides.
- C. Configure servers running in the VPC using the host-based 'route' commands to send all traffic through the platform to a scalable virtualized IDS/IPS.
- D. Configure each host with an agent that collects all network traffic and sends that traffic to the IDS/IPS platform for inspection.

Answer: C

NEW QUESTION 421

A web-startup runs its very successful social news application on Amazon EC2 with an Elastic Load Balancer, an Auto-Scaling group of Java/Tomcat application servers, and DynamoDB as data store. The main web-application best runs on m2.xlarge instances since it is highly memory-bound. Each new deployment requires semi-automated creation and testing of a new AM for the application servers which takes quite a while and is therefore only done once per week. Recently, a new chat feature has been implemented in Node.js and needs to be integrated in the architecture. First tests show that the new component is CPU bound. Because the company has some experience with using Chef, they decided to streamline the deployment process and use AWS Ops Works as an application life cycle tool to simplify management of the application and reduce the deployment cycles.

What configuration in AWS Ops Works is necessary to integrate the new chat module in the most cost-efficient and flexible way?

- A. Create one AWS OpsWorks stack, create one AWS OpsWorks layer, create one custom recipe.
- B. Create one AWS OpsWorks stack, create two AWS OpsWorks layers, create one custom recipe.

- C. Create two AWS OpsWorks stacks create two AWS Ops Works layers create one custom recipe
- D. Create two AWS OpsWorks stacks create two AWS Ops Works layers create two custom recipe

Answer: C

NEW QUESTION 424

Your firm has uploaded a large amount of aerial image data to S3. In the past, in your on-premises environment, you used a dedicated group of servers to process this data and used RabbitMQ an open source messaging system to get job information to the servers. Once processed the data would go to tape and be shipped offsite. Your manager told you to stay with the current design, and leverage AWS archival storage and messaging services to minimize cost. Which is correct?

- A. Use SNS for passing job messages use Cloud Watch alarms to terminate EC2 worker instances when they become idle
- B. Once data is processed, change the storage class of the S3 objects to Reduced Redundancy Storage.
- C. Setup Auto-Scaled workers triggered by queue depth that use spot instances to process messages in SNS. Once data is processed,
- D. Change the storage class of the S3 objects to Reduced Redundancy Storage
- E. Setup Auto-Scaled workers triggered by queue depth that use spot instances to process messages in SNS. Once data is processed, change the storage class of the S3 objects to Glacier.
- F. Use SNS to pass job messages use Cloud Watch alarms to terminate spot worker instances when they become idle
- G. Once data is processed, change the storage class of the S3 object to Glacier.

Answer: D

NEW QUESTION 426

What does Amazon S3 stand for?

- A. Simple Storage Solution.
- B. Storage Storage Storage (triple redundancy Storage).
- C. Storage Sewer Solution.
- D. Simple Storage Sewer

Answer: D

NEW QUESTION 430

Select the most correct

The device name /dev/sda1 (within Amazon EC2) is _

- A. Possible for EBS volumes
- B. Reserved for the root device
- C. Recommended for EBS volumes
- D. Recommended for instance store volumes

Answer: B

NEW QUESTION 434

If I write the below command, what does it do? `ec2-run ami-e3a5408a -n 20 -g appserver`

- A. Start twenty instances as members of appserver group.
- B. Creates 20 rules in the security group named appserver
- C. Terminate twenty instances as members of appserver group.
- D. Start 20 security groups

Answer: A

NEW QUESTION 435

Every user you create in the IAM system starts with _ _

- A. Partial permissions
- B. Full permissions
- C. No permissions

Answer: C

NEW QUESTION 436

Can I control if and when MySQL based RDS Instance is upgraded to new supported versions?

- A. No
- B. Only in VPC
- C. Yes

Answer: C

NEW QUESTION 437

When you view the block device mapping for your instance, you can see only the EBS volumes, not the instance store volumes.

- A. Depends on the instance type

- B. FALSE
- C. Depends on whether you use API call
- D. TRUE

Answer: D

NEW QUESTION 442

What are the initial settings of a user created security group?

- A. Allow all inbound traffic and Allow no outbound traffic
- B. All low no inbound traffic and All low no outbound traffic
- C. All low no inbound traffic and All low all outbound traffic
- D. Allow all inbound traffic and Allow all outbound traffic

Answer: C

NEW QUESTION 446

Can we attach an EBS volume to more than one EC2 instance at the same time?

- A. No
- B. Yes.
- C. Only EC2-optimized EBS volumes.
- D. Only in read mod

Answer: A

NEW QUESTION 449

What does the AWS Storage Gateway provide?

- A. It allows to integrate on-premises IT environments with Cloud Storage.
- B. A direct encrypted connection to Amazon S3.
- C. It's a backup solution that provides an on-premises Cloud storage.
- D. It provides an encrypted SSL endpoint for backups in the Cloud

Answer: A

NEW QUESTION 454

How many relational database engines does RDS currently support?

- A. Three: MySQL, Oracle and Microsoft SQL Server.
- B. Just two: MySQL and Oracle.
- C. Five: MySQL, PostgreSQL, MongoDB, Cassandra and SQLite.
- D. Just one: MySQL

Answer: A

NEW QUESTION 459

Fill in the blanks: The base URI for all requests for instance metadata is __ __

- A. <http://254.169.169.254/latest/>
- B. <http://169.169.254.254/latest/>
- C. <http://127.0.0.1/latest/>
- D. <http://169.254.169.254/latest/>

Answer: D

NEW QUESTION 463

What are the two permission types used by AWS?

- A. Resource-based and Product-based
- B. Product-based and Service-based
- C. Service-based
- D. User-based and Resource-based

Answer: D

NEW QUESTION 466

In the Amazon CloudWatch, which metric should I be checking to ensure that your DB Instance has enough free storage space?

- A. Free Storage
- B. Free Storage Space
- C. Free Storage Volume
- D. Free DB Storage Space

Answer: B

NEW QUESTION 467

Which is the default region in AWS?

- A. eu-west-1
- B. us-east-1
- C. us-east-2
- D. ap-southeast-1

Answer: B

NEW QUESTION 468

What are the Amazon EC2 API tools?

- A. They don't exist
- B. The Amazon EC2 AMI tools, instead, are used to manage permissions.
- C. Command-line tools to the Amazon EC2 web service.
- D. They are a set of graphical tools to manage EC2 instances.
- E. They don't exist
- F. The Amazon API tools are a client interface to Amazon Web Service

Answer: B

NEW QUESTION 470

EBS Snapshots occur _

- A. Asynchronously
- B. Synchronously
- C. Weekly

Answer: A

NEW QUESTION 472

Disabling automated backups _ disable the point-in-time recovery.

- A. if configured to can
- B. will never
- C. will

Answer: C

NEW QUESTION 473

Is creating a Read Replica of another Read Replica supported?

- A. Only in certain regions
- B. Only with MSSQL based RDS
- C. Only for Oracle RDS types
- D. No

Answer: D

NEW QUESTION 475

Can Amazon S3 uploads resume on failure or do they need to restart?

- A. Restart from beginning
- B. You can resume them, if you flag the "resume on failure" option before uploading.
- C. Resume on failure
- D. Depends on the file size

Answer: C

NEW QUESTION 479

What does the following command do with respect to the Amazon EC2 security groups? `ec2-revoke RevokeSecurityGroup Ingress`

- A. Removes one or more security groups from a rule.
- B. Removes one or more security groups from an Amazon EC2 instance.
- C. Removes one or more rules from a security group.
- D. Removes a security group from our account

Answer: C

NEW QUESTION 480

Can a 'user' be associated with multiple AWS accounts?

- A. No

B. Yes

Answer: A

NEW QUESTION 485

Will my standby RDS instance be in the same Availability Zone as my primary?

- A. Only for Oracle RDS types
- B. Yes
- C. Only if configured at launch
- D. No

Answer: D

NEW QUESTION 489

While launching an RDS DB instance, on which page I can select the Availability Zone?

- A. REVIEW
- B. DB INSTANCE DETAILS
- C. MANAGEMENT OPTIONS
- D. ADDITIONAL CONFIGURATION

Answer: D

NEW QUESTION 494

In the Launch Db Instance Wizard, where can I select the backup and maintenance options?

- A. Under DB INSTANCE DETAILS
- B. Under REVIEW
- C. Under MANAGEMENT OPTIONS
- D. Under ENGINE SELECTION

Answer: C

NEW QUESTION 498

By default, when an EBS volume is attached to a Windows instance, it may show up as any drive letter on the instance. You can change the settings of the _ Senrice to set the drive letters of the EBS volumes per your specifications.

- A. EBS Config Senrice
- B. AMI Config Senrice
- C. Ec2 Config Senrice
- D. Ec2-AMI Config Senrice

Answer: C

NEW QUESTION 499

What is Oracle SQL Developer?

- A. An AWS developer who is an expert in Amazon RDS using both the Oracle and SQL Server DB engines
- B. A graphical Java tool distributed without cost by Oracle.
- C. It is a variant of the SQL Server Management Studio designed by Microsoft to support Oracle DBMS functionalities
- D. A different DBMS released by Microsoft free of cost

Answer: B

NEW QUESTION 503

To view information about an Amazon EBS volume, open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>, click in the Navigation pane.

- A. EBS
- B. Describe
- C. Details
- D. Volumes

Answer: D

NEW QUESTION 506

Using Amazon IAM, can I give permission based on organizational groups?

- A. Yes but only in certain cases
- B. No
- C. Yes always

Answer: C

NEW QUESTION 511

While signing in REST/ Query requests, for additional security, you should transmit your requests using Secure Sockets Layer (SSL) by using

- A. HTIP
- B. Internet Protocol Security(IPsec)
- C. TLS (Transport Layer Security)
- D. HTIPS

Answer: D

NEW QUESTION 514

What happens to the I/O operations while you take a database snapshot?

- A. I/O operations to the database are suspended for a few minutes while the backup is in progress.
- B. I/O operations to the database are sent to a Replica (if available) for a few minutes while the backup is in progress.
- C. I/O operations will be functioning normally
- D. I/O operations to the database are suspended for an hour while the backup is in progress

Answer: A

NEW QUESTION 519

When running my DB Instance as a Multi-AZ deployment, can I use the standby for read or write operations?

- A. Yes
- B. Only with MSSQL based RDS
- C. Only for Oracle RDS instances
- D. No

Answer: D

NEW QUESTION 520

In the 'Detailed' monitoring data available for your Amazon EBS volumes, Provisioned IOPS volumes automatically send _ minute metrics to Amazon CloudWatch.

- A. 3
- B. 1
- C. 5
- D. 2

Answer: B

NEW QUESTION 523

What is the minimum charge for the data transferred between Amazon RDS and Amazon EC2 Instances in the same Availability Zone?

- A. USD 0.10 per GB
- B. No charge
- C. It is free.
- D. USD 0.02 per GB
- E. USD 0.01 per GB

Answer: B

NEW QUESTION 528

Which Amazon Storage behaves like raw, unformatted, external block devices that you can attach to your instances?

- A. None of these.
- B. Amazon Instance Storage
- C. Amazon EBS
- D. All of these

Answer: C

NEW QUESTION 531

Which Amazon service can I use to define a virtual network that closely resembles a traditional data center?

- A. Amazon VPC
- B. Amazon Service Bus
- C. Amazon EMR
- D. Amazon RDS

Answer: A

NEW QUESTION 536

If I modify a DB Instance or the DB parameter group associated with the instance, should I reboot the instance for the changes to take effect?

- A. No

B. Yes

Answer: B

NEW QUESTION 540

Can I test my DB Instance against a new version before upgrading?

- A. No
- B. Yes
- C. Only in VPC

Answer: B

NEW QUESTION 542

Fill in the blanks: "To ensure failover capabilities, consider using a _ for incoming traffic on a network interface".

- A. primary public IP
- B. secondary private I P
- C. secondary public I P
- D. add on secondary IP

Answer: B

NEW QUESTION 545

Can I encrypt connections between my application and my DB Instance using SSL?

- A. No
- B. Yes
- C. Only in VPC
- D. Only in certain regions

Answer: B

NEW QUESTION 546

What are the four levels of AWS Premium Support?

- A. Basic, Developer, Business, Enterprise
- B. Basic, Startup, Business, Enterprise
- C. Free, Bronze, Silver, Gold
- D. All support is free

Answer: A

NEW QUESTION 549

What can I access by visiting the URL: <http://status.aws.amazon.com/>?

- A. Amazon Cloud Watch
- B. Status of the Amazon RDS DB
- C. AWS Service Health Dashboard
- D. AWS Cloud Monitor

Answer: C

NEW QUESTION 552

Please select the Amazon EC2 resource which cannot be tagged.

- A. images (AMIs, kernels, RAM disks)
- B. Amazon EBS volumes
- C. Elastic IP addresses
- D. VPCs

Answer: C

NEW QUESTION 554

Because of the extensibility limitations of striped storage attached to Windows Server, Amazon RDS does not currently support increasing storage on a _ DB Instance.

- A. SQL Server
- B. MySQL
- C. Oracle

Answer: A

NEW QUESTION 557

How are the EBS snapshots saved on Amazon S3?

- A. Exponentially
- B. Incrementally
- C. EBS snapshots are not stored in the Amazon S3
- D. Decrementally

Answer: B

NEW QUESTION 559

Can I delete a snapshot of the root device of an EBS volume used by a registered AMI?

- A. Only via API
- B. Only via Console
- C. Yes
- D. No

Answer: C

NEW QUESTION 563

True or False: Without IAM, you cannot control the tasks a particular user or system can do and what AWS resources they might use.

- A. FALSE
- B. TRUE

Answer: A

NEW QUESTION 564

When automatic failover occurs, Amazon RDS will emit a DB Instance event to inform you that automatic failover occurred. You can use the to return information about events related to your DB Instance

- A. FetchFailure
- B. DescribeFailure
- C. DescribeEvents
- D. FetchEvents

Answer: C

NEW QUESTION 566

What is the default maximum number of MFA devices in use per AWS account (at the root account level)?

- A. 1
- B. 5
- C. 15
- D. 10

Answer: A

NEW QUESTION 570

Do the Amazon EBS volumes persist independently from the running life of an Amazon EC2 instance?

- A. Only if instructed to when created
- B. Yes
- C. No

Answer: B

NEW QUESTION 573

Please select the Amazon EC2 resource which can be tagged.

- A. key pairs
- B. Elastic IP addresses
- C. placement groups
- D. Amazon EBS snapshots

Answer: C

NEW QUESTION 575

Select the correct statement:

- A. You don't need to specify the resource identifier while stopping a resource
- B. You can terminate, stop, or delete a resource based solely on its tags
- C. You can't terminate, stop, or delete a resource based solely on its tags
- D. You don't need to specify the resource identifier while terminating a resource

Answer: C

NEW QUESTION 577

A group can contain many users. Can a user belong to multiple groups?

- A. Yes always
- B. No
- C. Yes but only if they are using two factor authentication
- D. Yes but only in VPC

Answer: A

NEW QUESTION 579

Is the encryption of connections between my application and my DB Instance using SSL for the MySQL server engines available?

- A. Yes
- B. Only in VPC
- C. Only in certain regions
- D. No

Answer: A

NEW QUESTION 582

Which AWS instance address has the following characteristics? : " If you stop an instance, its Elastic IP address is unmapped, and you must remap it when you restart the instance."

- A. Both A and B
- B. None of these
- C. VPC Addresses
- D. EC2 Addresses

Answer: A

NEW QUESTION 585

Multi-AZ deployment _ supported for Microsoft SQL Server DB Instances.

- A. is not currently
- B. is as of 2013
- C. is planned to be in 2014
- D. will never be

Answer: A

NEW QUESTION 587

Security groups act like a firewall at the instance level, whereas _ are an additional layer of security that act at the subnet level.

- A. DB Security Groups
- B. VPC Security Groups
- C. network ACLs

Answer: C

NEW QUESTION 591

Are you able to integrate a multi-factor token service with the AWS Platform?

- A. Yes, using the AWS multi-factor token devices to authenticate users on the AWS platform.
- B. No, you cannot integrate multi-factor token devices with the AWS platform.
- C. Yes, you can integrate private multi-factor token devices to authenticate users to the AWS platfor

Answer: A

NEW QUESTION 594

My Read Replica appears "stuck" after a Multi-AZ failover and is unable to obtain or apply updates from the source DB Instance. What do I do?

- A. You will need to delete the Read Replica and create a new one to replace it.
- B. You will need to disassociate the DB Engine and re associate it.
- C. The instance should be deployed to Single AZ and then moved to Multi- AZ once again
- D. You will need to delete the DB Instance and create a new one to replace i

Answer: A

NEW QUESTION 595

Does Route 53 support MX Records?

- A. Yes.
- B. It supports CNAME records, but not MX records.
- C. No
- D. Only Primary MX record
- E. Secondary MX records are not supported

Answer: A

NEW QUESTION 596

Which Amazon storage do you think is the best for my database-style applications that frequently encounter many random reads and writes across the dataset?

- A. None of these.
- B. Amazon Instance Storage
- C. Any of these
- D. Amazon EBS

Answer: D

NEW QUESTION 598

Select the correct set of steps for exposing the snapshot only to specific AWS accounts

- A. Select public for all the accounts and check mark those accounts with whom you want to expose the snapshots and click save.
- B. Select Private, enter the IDs of those AWS accounts, and click Save.
- C. Select Public, enter the IDs of those AWS accounts, and click Save.
- D. Select Public, mark the IDs of those AWS accounts as private, and click Save

Answer: C

NEW QUESTION 602

Is decreasing the storage size of a DB Instance permitted?

- A. Depends on the ROMS used
- B. Yes
- C. No

Answer: B

NEW QUESTION 605

In the context of MySQL, version numbers are organized as MySQL version = X.Y.Z. What does X denote here?

- A. release level
- B. minor version
- C. version number
- D. major version

Answer: D

NEW QUESTION 608

In the 'Detailed' monitoring data available for your Amazon EBS volumes, Provisioned IOPS volumes automatically send _ minute metrics to Amazon CloudWatch.

- A. 5
- B. 2
- C. 1
- D. 3

Answer: C

NEW QUESTION 611

Are you able to integrate a multi-factor token service with the AWS Platform?

- A. Yes, you can integrate private multi-factor token devices to authenticate users to the AWS platform.
- B. No, you cannot integrate multi-factor token devices with the AWS platform.
- C. Yes, using the AWS multi-factor token devices to authenticate users on the AWS platform.

Answer: C

NEW QUESTION 616

You can use _ and _ to help secure the instances in your VPC,

- A. security groups and multi-factor authentication
- B. security groups and 2-Factor authentication
- C. security groups and biometric authentication
- D. security groups and network ACLs

Answer: D

NEW QUESTION 620

What is the maximum response time for a Business level Premium Support case?

- A. 30 minutes
- B. You always get instant responses (within a few seconds).
- C. 10 minutes
- D. 1 hour

Answer: D

NEW QUESTION 624

What does Amazon ELB stand for?

- A. Elastic Linux Box.
- B. Encrypted Linux Box.
- C. Encrypted Load Balancing.
- D. Elastic Load Balancing

Answer: D

NEW QUESTION 626

What does Amazon Cloud Formation provide?

- A. None of these.
- B. The ability to setup Autoscaling for Amazon EC2 instances.
- C. A template to map network resources for Amazon Web Services.
- D. A templated resource creation for Amazon Web Service

Answer: D

NEW QUESTION 627

Is there a limit to the number of groups you can have?

- A. Yes for all users except root
- B. No
- C. Yes unless special permission granted
- D. Yes for all users

Answer: D

NEW QUESTION 628

Is there a method in the IAM system to allow or deny access to a specific instance?

- A. Only for VPC based instances
- B. Yes
- C. No

Answer: C

NEW QUESTION 629

What does Amazon SES stand for?

- A. Simple Elastic Server
- B. Simple Email Service
- C. Software Email Solution
- D. Software Enabled Server

Answer: B

NEW QUESTION 633

Amazon S3 doesn't automatically give a user who creates _ permission to perform other actions on that bucket or object.

- A. a file
- B. a bucket or object
- C. a bucket or file
- D. a object or file

Answer: B

NEW QUESTION 637

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AWS-Solution-Architect-Associate Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AWS-Solution-Architect-Associate Product From:

<https://www.2passeasy.com/dumps/AWS-Solution-Architect-Associate/>

Money Back Guarantee

AWS-Solution-Architect-Associate Practice Exam Features:

- * AWS-Solution-Architect-Associate Questions and Answers Updated Frequently
- * AWS-Solution-Architect-Associate Practice Questions Verified by Expert Senior Certified Staff
- * AWS-Solution-Architect-Associate Most Realistic Questions that Guarantee you a Pass on Your First Try
- * AWS-Solution-Architect-Associate Practice Test Questions in Multiple Choice Formats and Updates for 1 Year