AWS-Solution-Architect-Associate Dumps

Amazon AWS Certified Solutions Architect - Associate

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

NEW QUESTION 1
You are trying to launch an EC2 instance, however the instance seems to go into a terminated status immediately. What would probably not be a reason that this is happening?

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

Answer: C

Explanation:
Amazon EC2 provides a virtual computing environments, known as an instance.
After you launch an instance, AWS recommends that you check its status to confirm that it goes from the pending status to the running status, the not terminated status.
The following are a few reasons why an Amazon EBS-backed instance might immediately terminate: You've reached your volume limit.
The AM is missing a required part. The snapshot is corrupt. Reference:
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_InstanceStraightToTerminated.html

NEW QUESTION 2
In Amazon EC2 Container Service components, what is the name of a logical grouping of container instances on which you can place tasks?

A. A cluster
B. A container instance
C. A container
D. A task definition

Answer: A

Explanation:
Amazon ECS contains the following components:
A Cluster is a logical grouping of container instances that you can place tasks on.
A Container instance is an Amazon EC2 instance that is running the Amazon ECS agent and has been registered into a cluster.
A Task definition is a description of an application that contains one or more container definitions. A Scheduler is the method used for placing tasks on container instances.
A Service is an Amazon ECS service that allows you to run and maintain a specified number of instances of a task definition simultaneously.
A Task is an instantiation of a task definition that is running on a container instance. A Container is a Linux container that was created as part of a task.
Reference: http://docs.aws.amazon.com/AmazonECS/latest/developerguide/Welcome.html

NEW QUESTION 3
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 4
After you recommend Amazon Redshift to a client as an alternative solution to paying data warehouses to analyze his data, your client asks you to explain why you are recommending Redshift. Which of the following would be a reasonable response to his request?

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

Answer: D

Explanation:
Amazon Redshift delivers fast query performance by using columnar storage technology to improve I/O efficiency and parallelizing queries across multiple nodes.
Redshift uses standard PostgreSQL JDBC and ODBC drivers, allowing you to use a wide range of familiar SQL clients. Data load speed scales linearly with cluster size, with integrations to Amazon S3, Amazon DynamoDB, Amazon Elastic MapReduce,
Amazon Kinesis or any SSH-enabled host.
AWS recommends Amazon Redshift for customers who have a combination of needs, such as: High performance at scale as data and query complexity grows
Desire to prevent reporting and analytic processing from interfering with the performance of OLTP workloads
Large volumes of structured data to persist and query using standard SQL and existing BI tools Desire to the administrative burden of running one's own data
warehouse and dealing with setup, durability, monitoring, scaling and patching
Reference: https://aws.amazon.com/running_databases/#redshift_anchor

NEW QUESTION 5
After setting up a Virtual Private Cloud (VPC) network, a more experienced cloud engineer suggests that to achieve low network latency and high network throughput you should look into setting up a placement group. You know nothing about this, but begin to do some research about it and are especially curious about its limitations. Which of the below statements is wrong in describing the limitations of a placement group?

A. Although launching multiple instance types into a placement group is possible, this reduces the likelihood that the required capacity will be available for your launch to succeed.
B. A placement group can span multiple Availability Zones.
C. You can't move an existing instance into a placement group.
D. A placement group can span peered VPCs

Answer: B

Explanation:
A placement group is a logical grouping of instances within a single Availability Zone. Using placement groups enables applications to participate in a low-latency, 10 Gbps network. Placement groups are recommended for applications that benefit from low network latency, high network throughput, or both. To provide the lowest latency, and the highest packet-per-second network performance for your placement group, choose an instance type that supports enhanced networking. Placement groups have the following limitations:
The name you specify for a placement group a name must be unique within your AWS account. A placement group can't span multiple Availability Zones. Although launching multiple instance types into a placement group is possible, this reduces the likelihood that the required capacity will be available for your launch to succeed. We recommend using the same instance type for all instances in a placement group.
You can't merge placement groups. Instead, you must terminate the instances in one placement group, and then relaunch those instances into the other placement group.
A placement group can span peered VPCs; however, you will not get full-bisection bandwidth between instances in peered VPCs. For more information about VPC peering connections, see VPC Peering in the Amazon VPC User Guide.
You can't move an existing instance into a placement group. You can create an AM from your existing instance, and then launch a new instance from the AMI into a placement group.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html

NEW QUESTION 6
What is a placement group in Amazon EC2?

A. It is a group of EC2 instances within a single Availability Zone.
B. It the edge location of your web content.
C. It is the AWS region where you run the EC2 instance of your web content.
D. It is a group used to span multiple Availability Zone

Answer: A

Explanation:
A placement group is a logical grouping of instances within a single Availability Zone. Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html

NEW QUESTION 7
You are migrating an internal sewer 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

Your supervisor has asked you to build a simple file synchronization service for your department. He doesn't want to spend too much money and he wants to be notified of any changes to files by email. What do you think would be the best Amazon service to use for the email solution?

A. Amazon SES
B. Amazon CloudSearch
C. Amazon SWF
D. Amazon AppStream

Answer: A

Explanation:
File change notifications can be sent via email to users following the resource with Amazon Simple Email Service (Amazon SES), an easy-to-use, cost-effective email solution.
Reference: http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_filesync_08.pdf

NEW QUESTION 10
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/\NorkingWithItems.html#Working WithItems.AtomicCounters

NEW QUESTION 10
Amazon EC2 provides a . It is an HTTP or HTTPS request that uses the HTTP verbs GET or POST.

A. web database
B. .net framework
C. Query API
D. C library

Answer: C

Explanation:
Amazon EC2 provides a Query API. These requests are HTTP or HTTPS requests that use the HTTP verbs GET or POST and a Query parameter named Action.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/APIReference/making-api-requests.html

NEW QUESTION 13
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 15
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 20

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 23
You need to migrate a large amount of data into the cloud that you have stored on a hard disk and you decide that the best way to accomplish this is with AWS Import/Export and you mail the hard disk to AWS. Which of the following statements is incorrect in regards to AWS Import/Export?

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

Answer: C

Explanation:
AWS Import/Export supports: Import to Amazon S3
Export from Amazon S3 Import to Amazon EBS Import to Amazon Glacier
AWS Import/Export does not currently support export from Amazon EBS or Amazon Glacier. Reference:
https://docs.aws.amazon.com/AWSImportExport/latest/DG/whatisdisk.html

NEW QUESTION 27
In Amazon EC2, partial instance-hours are billed .

A. per second used in the hour
B. per minute used
C. by combining partial segments into full hours
D. as full hours

Answer: D

Explanation:
Partial instance-hours are billed to the next hour. Reference: http://aws.amazon.com/ec2/faqs/

NEW QUESTION 32
In EC2, what happens to the data in an instance store if an instance reboots (either intentionally or unintentionally)?

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

Answer: B

Explanation:
The data in an instance store persists only during the lifetime of its associated instance. If an instance reboots (intentionally or unintentionally), data in the instance store persists. However, data on instance store volumes is lost under the following circumstances.
Failure of an underlying drive
Stopping an Amazon EBS-backed instance Terminating an instance
Reference: http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/InstanceStorage.html

NEW QUESTION 37
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 subne

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 42
While using the EC2 GET requests as URLs, the is the URL that serves as the entry point for the web service.

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

Answer: B

Explanation:
The endpoint is the URL that serves as the entry point for the web service.
Reference: http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/using-query-api.htmI

NEW QUESTION 45
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_co|umnar_storage_disk_mem_mgmnt.html

NEW QUESTION 47
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/DeveIoperGuide/GettingStartedTutoriaI.html

NEW QUESTION 49
You have been given a scope to deploy some AWS infrastructure for a large organisation. The requirements are that you will have a lot of EC2 instances but may need to add more when the average utilization of your Amazon EC2 fiIeet is high and conversely remove them when CPU utilization is low. Which AWS services would be best to use to accomplish this?

A. Auto Scaling, Amazon CIoudWatch and AWS Elastic Beanstalk
B. Auto Scaling, Amazon CIoudWatch and Elastic Load Balancing.
C. Amazon CIoudFront, Amazon CIoudWatch and Elastic Load Balancing.
D. AWS Elastic Beanstalk , Amazon CIoudWatch and Elastic Load Balancin

Answer: B

Explanation:
Auto Scaling enables you to follow the demand curve for your applications closely, reducing the need to manually provision Amazon EC2 capacity in advance. For example, you can set a condition to add new
Amazon EC2 instances in increments to the Auto Scaling group when the average utilization of your Amazon EC2 fiIeet is high; and similarly, you can set a condition to remove instances in the same increments when CPU utilization is low. If you have predictable load changes, you can set a schedule through Auto Scaling to plan your scaling actMties. You can use Amazon CIoudWatch to send alarms to trigger scaling actMties and Elastic Load Balancing to help distribute traffic to your instances within Auto Scaling groups. Auto Scaling enables you to run your Amazon EC2 fiIeet at optimal utilization. Reference:
http://aws.amazon.com/autoscaling/

NEW QUESTION 52
You've created your first load balancer and have registered your EC2 instances with the load balancer. Elastic Load Balancing routinely performs health checks on

all the registered EC2 instances and automatically distributes all incoming requests to the DNS name of your load balancer across your registered, healthy EC2 instances. By default, the load balancer uses the _ protocol for checking the health of your instances.

A. HTTPS
B. HTTP
C. ICMP
D. IPv6

Answer: B

Explanation:
In Elastic Load Balancing a health configuration uses information such as protocol, ping port, ping path (URL), response timeout period, and health check interval to determine the health state of the instances registered with the load balancer.
Currently, HTTP on port 80 is the default health check. Reference:
http://docs.aws.amazon.com/E|asticLoadBaIancing/latest/DeveIoperGuide/TerminoIogyandKeyConcepts. html

NEW QUESTION 54
A major finance organisation has engaged your company to set up a large data mining application. Using AWS you decide the best service for this is Amazon Elastic MapReduce(EMR) which you know uses Hadoop. Which of the following statements best describes Hadoop?

A. Hadoop is 3rd Party software which can be installed using AMI
B. Hadoop is an open source python web framework
C. Hadoop is an open source Java software framework
D. Hadoop is an open source javascript framework

Answer: C

Explanation:
Amazon EMR uses Apache Hadoop as its distributed data processing engine.
Hadoop is an open source, Java software framework that supports data-intensive distributed applications running on large clusters of commodity hardware.
Hadoop implements a programming model named "MapReduce," where the data is dMded into many small fragments of work, each of which may be executed on any node in the cluster.
This framework has been widely used by developers, enterprises and startups and has proven to be a reliable software platform for processing up to petabytes of data on clusters of thousands of commodity machines.
Reference: http://aws.amazon.com/elasticmapreduce/faqs/

NEW QUESTION 55
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_I2_as

NEW QUESTION 56
As AWS grows, most of your clients' main concerns seem to be about security, especially when all of their competitors also seem to be using AWS. One of your clients asks you whether having a competitor who hosts their EC2 instances on the same physical host would make it easier for the competitor to hack into the client's data. Which of the following statements would be the best choice to put your client's mind at rest?

A. Different instances running on the same physical machine are isolated from each other via a 256-bit Advanced Encryption Standard (AES-256).
B. Different instances running on the same physical machine are isolated from each other via the Xen hypervisor and via a 256-bit Advanced Encryption Standard (AES-256).
C. Different instances running on the same physical machine are isolated from each other via the Xen hypervisor.
D. Different instances running on the same physical machine are isolated from each other via IAM permissions.

Answer: C

Explanation:
Amazon Elastic Compute Cloud (EC2) is a key component in Amazon's Infrastructure as a Service (IaaS), providing resizable computing capacity using server instances in AWS's data centers. Amazon EC2 is designed to make web-scale computing easier by enabling you to obtain and configure capacity with minimal friction.
You create and launch instances, which are collections of platform hardware and software. Different instances running on the same physical machine are isolated from each other via the Xen hypervisor.
Amazon is active in the Xen community, which provides awareness of the latest developments. In addition, the AWS firewall resides within the hypervisor layer, between the physical network interface and the instance's virtual interface. All packets must pass through this layer, thus an instance's neighbors have no more access to that instance than any other host on the Internet and can be treated as if they are on
separate physical hosts. The physical RAM is separated using similar mechanisms.
Reference: http://d0.awsstatic.com/whitepapers/Security/AWS%20Security%20Whitepaper.pdf

NEW QUESTION 59
In Amazon RDS, security groups are ideally used to:

A. Define maintenance period for database engines
B. Launch Amazon RDS instances in a subnet
C. Create, describe, modify, and delete DB instances
D. Control what IP addresses or EC2 instances can connect to your databases on a DB instance

Answer: D

Explanation:
In Amazon RDS, security groups are used to control what IP addresses or EC2 instances can connect to your databases on a DB instance.
When you first create a DB instance, its firewall prevents any database access except through rules specified by an associated security group.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/UsingWithRDS.htmI

NEW QUESTION 61
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 SimpIeDB
C. N
D. Not presently
E. Ye
F. Amazon Redshift
G. Ye
H. Your choice of relational AMIs 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 65
In Amazon EC2, while sharing an Amazon EBS snapshot, can the snapshots with AWS IV|arketpIace product codes be public?

A. Yes, but only for US-based providers.
B. Yes, they can be public.
C. No, they cannot be made public.
D. Yes, they are automatically made public by the syste

Answer: C

Explanation:
Snapshots with AWS Marketplace product codes can't be made public. Reference:
http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/ebs-modifying-snapshot-permissions.ht ml

NEW QUESTION 69
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 70
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 wvvw.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/DeveIoperGuide/AboutHostedZones.html

NEW QUESTION 72
A user has launched 10 EC2 instances inside a placement group. Which of the below mentioned statements is true with respect to the placement group?

A. All instances must be in the same AZ
B. All instances can be across multiple regions
C. The placement group cannot have more than 5 instances
D. All instances must be in the same region

Answer: A

Explanation:
A placement group is a logical grouping of EC2 instances within a single Availability Zone. Using placement groups enables applications to participate in a low-latency, 10 Gbps network. Placement groups are recommended for applications that benefit from low network latency, high network throughput or both.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html

NEW QUESTION 73
An organization has developed a mobile application which allows end users to capture a photo on their mobile device, and store it inside an application. The application internally uploads the data to AWS S3. The organization wants each user to be able to directly upload data to S3 using their Google ID. How will the mobile app allow this?

A. Use the AWS Web identity federation for mobile applications, and use it to generate temporary security credentials for each user.
B. It is not possible to connect to AWS S3 with a Google ID.
C. Create an IAM user every time a user registers with their Google ID and use IAM to upload files to S3.
D. Create a bucket policy with a condition which allows everyone to upload if the login ID has a Google part to it.

Answer: A

Explanation:
For Amazon Web Services, the Web identity federation allows you to create cloud-backed mobile apps that use public identity providers, such as login with Facebook, Google, or Amazon. It will create temporary security credentials for each user, which will be authenticated by the AWS services, such as S3.
Reference: http://docs.aws.amazon.com/STS/latest/UsingSTS/CreatingWIF.htmI

NEW QUESTION 76
You have been doing a lot of testing of your VPC Network by deliberately failing EC2 instances to test whether instances are failing over properly. Your customer who will be paying the AWS bill for all this asks you if he being charged for all these instances. You try to explain to him how the billing works on EC2 instances to the best of your knowledge. What would be an appropriate response to give to the customer
in regards to this?

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

Answer: C

Explanation:
Billing commences when Amazon EC2 initiates the boot sequence of an AM instance. Billing ends when the instance shuts down, which could occur through a web services command, by running "shutdown -h", or through instance failure.
Reference: http://aws.amazon.com/ec2/faqs/#BiIIing

NEW QUESTION 79
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 ntot 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 81
Once again your customers are concerned about the security of their sensitive data and with their latest enquiry ask about what happens to old storage devices on AWS. What would be the best answer to this QUESTION ?

A. AWS reformats the disks and uses them again.

B. AWS uses the techniques detailed in DoD 5220.22-M to destroy data as part of the decommissioning process.
C. AWS uses their own proprietary software to destroy data as part of the decommissioning process.
D. AWS uses a 3rd party security organization to destroy data as part of the decommissioning proces

Answer: B

Explanation:
When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized indMduals.
AWS uses the techniques detailed in DoD 5220.22-M ("Nationa| Industrial Security Program Operating ManuaI ") or NIST 800-88 ("Guidelines for Media Sanitization") to destroy data as part of the decommissioning process.
All decommissioned magnetic storage devices are degaussed and physically destroyed in accordance with industry-standard practices.
Reference: http://d0.awsstatic.com/whitepapers/Security/AWS%20Security%20Whitepaper.pdf

NEW QUESTION 85
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 receMng 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 indMual 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 90
You need to change some settings on Amazon Relational Database Service but you do not want the database to reboot immediately which you know might happen depending on the setting that you change. Which of the following will cause an immediate DB instance reboot to occur?

A. You change storage type from standard to PIOPS, and Apply Immediately is set to true.
B. You change the DB instance class, and Apply Immediately is set to false.
C. You change a static parameter in a DB parameter group.
D. You change the backup retention period for a DB instance from 0 to a nonzero value or from a nonzero value to 0, and Apply Immediately is set to false.

Answer: A

Explanation:
A DB instance outage can occur when a DB instance is rebooted, when the DB instance is put into a state that prevents access to it, and when the database is restarted. A reboot can occur when you manually reboot your DB instance or when you change a DB instance setting that requires a reboot before it can take effect.
A DB instance reboot occurs immediately when one of the following occurs:
You change the backup retention period for a DB instance from 0 to a nonzero value or from a nonzero value to 0 and set Apply Immediately to true.
You change the DB instance class, and Apply Immediately is set to true.
You change storage type from standard to PIOPS, and Apply Immediately is set to true.
A DB instance reboot occurs during the maintenance window when one of the following occurs:
You change the backup retention period for a DB instance from 0 to a nonzero value or from a nonzero value to 0, and Apply Immediately is set to false.
You change the DB instance class, and Apply Immediately is set to false. Reference:
http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Troub|eshooting.htm|#CHAP_TroubI eshooting.Security

NEW QUESTION 93
Amazon RDS provides high availability and failover support for DB instances using .

A. customized deployments
B. Appstream customizations
C. log events
D. MuIti-AZ deployments

Answer: D

Explanation:
Amazon RDS provides high availability and failover support for DB instances using MuIti-AZ deployments. MuIti-AZ deployments for Oracle, PostgreSQL, MySQL, and MariaDB DB instances use Amazon technology, while SQL Server DB instances use SQL Server Mrroring.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.IV|u|tiAZ.htmI

NEW QUESTION 96
A major customer has asked you to set up his AWS infrastructure so that it will be easy to recover in the case of a disaster of some sort. Which of the following is important when thinking about being able to quickly launch resources in AWS to ensure business continuity in case of a disaster?

A. Create and maintain AM|s of key sewers where fast recovery is required.
B. Regularly run your sewers, test them, and apply any software updates and configuration changes.
C. All items listed here are important when thinking about disaster recovery.
D. Ensure that you have all supporting custom software packages available in AW

Answer: C

Explanation:
In the event of a disaster to your AWS infrastructure you should be able to quickly launch resources in Amazon Web Services (AWS) to ensure business continuity.
The following are some key steps you should have in place for preparation:
1. Set up Amazon EC2 instances to replicate or mirror data.
2. Ensure that you have all supporting custom software packages available in AWS.
3. Create and maintain AMIs of key servers where fast recovery is required.
4. Regularly run these servers, test them, and apply any software updates and configuration changes.
5. Consider automating the provisioning of AWS resources.
Reference: http://d36cz9buwru1tt.cloudfront.net/AWS_Disaster_Recovery.pdf

NEW QUESTION 98
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 102
What would be the best way to retrieve the public IP address of your EC2 instance using the CLI?

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

Answer: D

Explanation:
To determine your instance's public IP address from within the instance, you can use instance metadata. Use the following command to access the public IP address: For Linux use, $ curl
http://169.254.169.254/latest/meta-data/public-ipv4, and for Windows use, $ wget http://169.254.169.254/latest/meta-data/public-ipv4.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-instance-addressing.htm|

NEW QUESTION 107
A user needs to run a batch process which runs for 10 minutes. This will only be run once, or at maximum twice, in the next month, so the processes will be temporary only. The process needs 15 X-Large instances. The process downloads the code from S3 on each instance when it is launched, and then generates a temporary log file. Once the instance is terminated, all the data will be lost. Which of the below mentioned pricing models should the user choose in this case?

A. Spot instance.
B. Reserved instance.
C. On-demand instance.
D. EBS optimized instanc

Answer: A

Explanation:
In Amazon Web Services, the spot instance is useful when the user wants to run a process temporarily. The spot instance can terminate the instance if the other user outbids the existing bid. In this case all storage is temporary and the data is not required to be persistent. Thus, the spot instance is a good option to save money.
Reference: http://aws.amazon.com/ec2/purchasing-options/spot-instances/

NEW QUESTION 109
Which of the following is NOT a characteristic of Amazon Elastic Compute Cloud (Amazon EC2)?

A. It can be used to launch as many or as few virtual servers as you need.
B. It increases the need to forecast traffic by providing dynamic IP addresses for static cloud computing.
C. It eliminates your need to invest in hardware up front, so you can develop and deploy applications faster.
D. It offers scalable computing capacity in the Amazon Web Services (AWS) clou

Answer: B

Explanation:
Amazon Elastic Compute Cloud (Amazon EC2) provides scalable computing capacity in the Amazon Web Services (AWS) cloud. Using Amazon EC2 eliminates your need to invest in hardware up front, so you can develop and deploy applications faster. You can use Amazon EC2 to launch as many or as few virtual servers as you need, configure security and networking, and manage storage. Amazon EC2 enables you
to scale up or down to handle changes in requirements or spikes in popularity, reducing your need to forecast traffic.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/concepts.html

NEW QUESTION 113
An existing client comes to you and says that he has heard that launching instances into a VPC (virtual private cloud) is a better strategy than launching instances into a EC2-classic which he knows is what you currently do. You suspect that he is correct and he has asked you to do some research about this and get back to him. Which of the following statements is true in regards to what ability launching your instances into a VPC instead of EC2-Classic gives you?

A. All of the things listed here.
B. Change security group membership for your instances while they're running
C. Assign static private IP addresses to your instances that persist across starts and stops
D. Define network interfaces, and attach one or more network interfaces to your instances

Answer: A

Explanation:
By launching your instances into a VPC instead of EC2-Classic, you gain the ability to: Assign static private IP addresses to your instances that persist across starts and stops Assign multiple IP addresses to your instances
Define network interfaces, and attach one or more network interfaces to your instances Change security group membership for your instances while they're running
Control the outbound traffic from your instances (egress filtering) in addition to controlling the inbound traffic to them (ingress filtering)
Add an additional layer of access control to your instances in the form of network access control lists (ACL)
Run your instances on single-tenant hardware
Reference: http://media.amazonwebservices.com/AWS_Cloud_Best_Practices.pdf

NEW QUESTION 118
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 CIRD format.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html

NEW QUESTION 121
You have been using T2 instances as your CPU requirements have not been that intensive. However you now start to think about larger instance types and start looking at M and IV|3 instances. You are a little confused as to the differences between them as they both seem to have the same ratio of CPU and memory. Which statement below is incorrect as to why you would use one over the other?

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

Answer: B

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

NEW QUESTION 122
You have set up an Elastic Load Balancer (ELB) with the usual default settings, which route each request independently to the application instance with the smallest load. However, someone has asked you to bind a user's session to a specific application instance so as to ensure that all requests coming from the user during the session will be sent to the same application instance. AWS has a feature to do this. What is it called?

A. Connection draining
B. Proxy protocol
C. Tagging
D. Sticky session

Answer: D

Explanation:
An Elastic Load Balancer(ELB) by default, routes each request independently to the application instance
with the smallest load. However, you can use the sticky session feature (also known as session affinity), which enables the load balancer to bind a user's session to a specific application instance. This ensures that all requests coming from the user during the session will be sent to the same application instance. The key to managing the sticky session is determining how long your load balancer should consistently route the user's request to the same application instance. If your application has its own session cookie, then you can set Elastic Load Balancing to create the session cookie to follow the duration specified by the application's session cookie. If your application does not have its own session cookie, then you can set Elastic Load Balancing to create a session cookie by specifying your own stickiness duration. You can associate stickiness duration for only HTTP/HTTPS load balancer listeners.
An application instance must always receive and send two cookies: A cookie that defines the stickiness duration and a special Elastic Load Balancing cookie named AWSELB, that has the mapping to the application instance.
Reference: http://docs.aws.amazon.com/E|asticLoadBaIancing/latest/DeveIoperGuide/TerminoIogyandKeyConcepts. htmI#session-stickiness

NEW QUESTION 125
A user wants to achieve High Availability with PostgreSQL DB. Which of the below mentioned functionalities helps achieve HA?

A. Mu|ti 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 128
After setting up several database instances in Amazon Relational Database Service (Amazon RDS) you decide that you need to track the performance and health of your databases. How can you do this?

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

Answer: C

Explanation:
Amazon Relational Database Service (Amazon RDS) is a web service that makes it easier to set up, operate, and scale a relational database in the cloud. It provides cost-efficient, resizeable capacity for an industry-standard relational database and manages common database administration tasks.
There are several ways you can track the performance and health of a database or a DB instance. You can:
Use the free Amazon CIoudWatch service to monitor the performance and health of a DB instance. Subscribe to Amazon RDS events to be notified when changes occur with a DB instance, DB snapshot, DB parameter group, or DB security group.
View, download, or watch database log files using the Amazon RDS console or Amazon RDS APIs. You can also query some database log files that are loaded into database tables.
Use the AWS CIoudTraiI service to record AWS calls made by your AWS account. The calls are recorded in log files and stored in an Amazon S3 bucket.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Monitoring.htmI

NEW QUESTION 132
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 devic

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 137
Which of the following strategies can be used to control access to your Amazon EC2 instances?

A. DB security groups
B. IAM policies
C. None of these
D. EC2 security groups

Answer: D

Explanation:

IAM policies allow you to specify what actions your IAM users are allowed to perform against your EC2 Instances. However, when it comes to access control, security groups are what you need in order to define and control the way you want your instances to be accessed, and whether or not certain kind of communications are allowed or not.
Reference: http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/UsingIAM.html

NEW QUESTION 139
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. AttachVoIume
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 141
Do you need to shutdown your EC2 instance when you create a snapshot of EBS volumes that serve as root devices?

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

Answer: B

Explanation:
Yes, to create a snapshot for Amazon EBS volumes that serve as root devices, you should stop the instance before taking the snapshot.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-creating-snapshot.html

NEW QUESTION 144
A user is running a batch process which runs for 1 hour every day. Which of the below mentioned options is the right instance type and costing model in this case if the user performs the same task for the whole year?

A. EBS backed instance with on-demand instance pricing.
B. EBS backed instance with heavy utilized reserved instance pricing.
C. EBS backed instance with low utilized reserved instance pricing.
D. Instance store backed instance with spot instance pricin

Answer: A

Explanation:
For Amazon Web Services, the reserved instance helps the user save money if the user is going to run the same instance for a longer period. Generally if the user uses the instances around 30-40% annually it is recommended to use RI. Here as the instance runs only for 1 hour daily it is not recommended to have RI as it will be costlier. The user should use on-demand with EBS in this case.
Reference: http://aws.amazon.com/ec2/purchasing-options/reserved-instances/

NEW QUESTION 146
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 150
You have a number of image files to encode. In an Amazon SQS worker queue, you create an Amazon SQS message for each file specifying the command (jpeg-encode) and the location of the file in Amazon S3. Which of the following statements best describes the functionality of Amazon SQS?

A. Amazon SQS is a distributed queuing system that is optimized for horizontal scalability, not for single-threaded sending or receMng speeds.
B. Amazon SQS is for single-threaded sending or receMng speeds.
C. Amazon SQS is a non-distributed queuing system.
D. Amazon SQS is a distributed queuing system that is optimized for vertical scalability and for single-threaded sending or receMng speeds.

Answer: A

Explanation:

Amazon SQS is a distributed queuing system that is optimized for horizontal scalability, not for
single-threaded sending or receMng speeds. A single client can send or receive Amazon SQS messages at a rate of about 5 to 50 messages per second. Higher receive performance can be achieved by requesting multiple messages (up to 10) in a single call. It may take several seconds before a message that has been to a queue is available to be received.
Reference: http://media.amazonwebservices.com/AWS_Storage_Options.pdf

NEW QUESTION 154
Select the correct statement: Within Amazon EC2, when using Linux instances, the device name
/dev/sda1 is .

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

Answer: D

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

NEW QUESTION 157
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 160
A user has set up the CloudWatch alarm on the CPU utilization metric at 50%, with a time interval of 5 minutes and 10 periods to monitor. What will be the state of the alarm at the end of 90 minutes, if the CPU utilization is constant at 80%?

A. ALERT
B. ALARM
C. OK
D. INSUFFICIENT_DATA

Answer: B

Explanation:
In this case the alarm watches a metric every 5 minutes for 10 intervals. Thus, it needs at least 50 minutes to come to the "OK" state.
Till then it will be in the |NSUFFUCIENT_DATA state.
Since 90 minutes have passed and CPU utilization is at 80% constant, the state of alarm will be "ALARNI". Reference:
http://docs.aws.amazon.com/AmazonCIoudWatch/latest/DeveloperGuide/AlarmThatSendsEmaiI.html

NEW QUESTION 163
A user comes to you and wants access to Amazon CloudWatch but only wants to monitor a specific LoadBaIancer. Is it possible to give him access to a specific set of instances or a specific LoadBaIancer?

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 LoadBaIancer. 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/AmazonC|oudWatch/latest/DeveloperGuide/UsingIAM.html

NEW QUESTION 168
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 173
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 176
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 C|oudTrai| 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 accoun

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 IV|anagement 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 177
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 181
You are configuring a new VPC for one of your clients for a cloud migration project, and only a public VPN will be in place. After you created your VPC, you created a new subnet, a new internet gateway, and attached your internet gateway to your VPC. When you launched your first instance into your VPC, you realized that you aren't able to connect to the instance, even if it is configured with an elastic IP. What should be done to access the instance?

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

Answer: A

Explanation:
All traffic should be routed via Internet Gateway. So, a route should be created with 0.0.0.0/0 as a source, and your Internet Gateway as your target.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario1.htmI

NEW QUESTION 183
A user is currently building a website which will require a large number of instances in six months, when a demonstration of the new site will be given upon launch. Which of the below mentioned options allows the user to procure the resources beforehand so that they need not worry about infrastructure availability during the demonstration?

A. Procure all the instances as reserved instances beforehand.
B. Launch all the instances as part of the cluster group to ensure resource availability.
C. Pre-warm all the instances one month prior to ensure resource availability.
D. Ask AWS now to procure the dedicated instances in 6 month

Answer: A

Explanation:
Amazon Web Services has massive hardware resources at its data centers, but they are finite. The best way for users to maximize their access to these resources is by reserving a portion of the computing capacity that they require. This can be done through reserved instances. With reserved instances, the user literally reserves the computing capacity in the Amazon Web Services cloud.
Reference: http://media.amazonwebservices.com/AWS_Building_Fault_To|erant_Applications.pdf

NEW QUESTION 187
You are setting up some EBS volumes for a customer who has requested a setup which includes a RAID (redundant array of inexpensive disks). AWS has some recommendations for RAID setups. Which RAID setup is not recommended for Amazon EBS?

A. RAID 5 only
B. RAID 5 and RAID 6
C. RAID 1 only
D. RAID 1 and RAID 6

Answer: B

Explanation:
With Amazon EBS, you can use any of the standard RAID configurations that you can use with a traditional bare metal server, as long as that particular RAID configuration is supported by the operating system for your instance. This is because all RAID is accomplished at the software level. For greater I/O performance than you can achieve with a single volume, RAID 0 can stripe multiple volumes together; for on-instance redundancy, RAID 1 can mirror two volumes together. RAID 5 and RAID 6 are not recommended for Amazon EBS because the parity write operations of these RAID modes consume some of the IOPS available to your volumes.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/raid-config.html

NEW QUESTION 189
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/swh/website-hosting-intro.html

NEW QUESTION 190
What is the network performance offered by the c4.8xIarge instance in Amazon EC2?

A. 20 Gigabit
B. 10 Gigabit
C. Very High but variable
D. 5 Gigabit

Answer: B

Explanation:
Networking performance offered by the c4.8xIarge instance is 10 Gigabit. Reference: http://aws.amazon.com/ec2/instance-types/

NEW QUESTION 194
How long does an AWS free usage tier EC2 last for?

A. Forever
B. 12 Months upon signup
C. 1 Month upon signup
D. 6 Months upon signup

Answer: B

Explanation:
The AWS free usage tier will expire 12 months from the date you sign up. When your free usage expires or if your application use exceeds the free usage tiers, you simply pay the standard, pay-as-you-go service rates.
Reference: http://aws.amazon.com/free/faqs/

NEW QUESTION 197
You have been setting up an Amazon Virtual Private Cloud (Amazon VPC) for your company, including setting up subnets. Security is a concern, and you are not sure which is the best security practice for securing subnets in your VPC. Which statement below is correct in describing the protection of AWS resources in each subnet?

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

Answer: A

Explanation:
A subnet is a range of IP addresses in your VPC. You can launch AWS resources into a subnet that you select. Use a public subnet for resources that must be connected to the Internet, and a private subnet for resources that won't be connected to the Internet.
To protect the AWS resources in each subnet, you can use multiple layers of security, including security groups and network access control lists (ACL).
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_|ntroduction.htmI

NEW QUESTION 202
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 syste

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 206
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 210
You have just finshed setting up an advertisement server in which one of the obvious choices for a service was Amazon Elastic Map Reduce(EMR) and are now troubleshooting some weird cluster states that you are seeing. Which of the below is not an Amazon EMR cluster state?

A. STARTING
B. STOPPED
C. RUNNING
D. WAITING

Answer: B

Explanation:
Amazon Elastic Map Reduce (EMR) is a web service that enables businesses, researchers, data analysts, and developers to easily and cost-effectively process vast amounts of data.

Amazon EMR historically referred to an Amazon EMR cluster (and all processing steps assigned to it) as a "cluster". Every cluster has a unique identifier that starts with "j-".

The different cluster states of an Amazon EMR cluster are listed below. STARTING — The cluster provisions, starts, and configures EC2 instances.

BOOTSTRAPPING — Bootstrap actions are being executed on the cluster. RUNNING — A step for the cluster is currently being run.

WAITING — The cluster is currently active, but has no steps to run. TERMINATING - The cluster is in the process of shutting down. TERMINATED - The cluster was shut down without error. TERMINATED_W|TH_ERRORS - The cluster was shut down with errors.

Reference: https://aws.amazon.com/elasticmapreduce/faqs/

NEW QUESTION 214

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 CloudHSNI 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 216

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 C|oudTrai| in the AWS Management Console.

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

NEW QUESTION 217

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 221

After moving an E-Commerce website for a client from a dedicated server to AWS you have also set up auto scaling to perform health checks on the instances in your group and replace instances that fail these checks. Your client has come to you with his own health check system that he wants you to use as it has proved to be very useful prior to his site running on AWS. What do you think would be an appropriate response to this given all that you know about auto scaling?

A. It is not possible to implement your own health check syste
B. You need to use AWSs health check system.
C. It is not possible to implement your own health check system due to compatibility issues.
D. It is possible to implement your own health check system and then send the instance's health information directly from your system to Cloud Watch.
E. It is possible to implement your own health check system and then send the instance's health information directly from your system to Cloud Watch but only in the US East (
F. Virginia) region.

Answer: C

Explanation:

Auto Scaling periodically performs health checks on the instances in your group and replaces instances that fail these checks. By default, these health checks use the results of EC2 instance status checks to determine the health of an instance. If you use a load balancer with your Auto Scaling group, you can optionally choose to include the results of Elastic Load Balancing health checks.

Auto Scaling marks an instance unhealthy if the calls to the Amazon EC2 action DescribeInstanceStatus returns any other state other than running, the system status shows impaired, or the calls to Elastic Load Balancing action DescribeInstanceHealth returns OutOfService in the instance state field.
After an instance is marked unhealthy because of an Amazon EC2 or Elastic Load Balancing health check, it is scheduled for replacement.
You can customize the health check conducted by your Auto Scaling group by specifying additional checks or by having your own health check system and then sending the instance's health information directly from your system to Auto Scaling.
Reference: http://docs.aws.amazon.com/AutoScaling/latest/Deve|operGuide/healthcheck.html

NEW QUESTION 223
You havejust discovered that you can upload your objects to Amazon S3 using Multipart Upload API. You start to test it out but are unsure of the benefits that it would provide. Which of the following is not a benefit of using multipart uploads?

A. You can begin an upload before you know the final object size.
B. Quick recovery from any network issues.
C. Pause and resume object uploads.
D. It's more secure than normal uploa

Answer: D

Explanation:
Multipart upload in Amazon S3 allows you to upload a single object as a set of parts. Each part is a contiguous portion ofthe object's data. You can upload these object parts independently and in any order.
If transmission of any part fails, you can re-transmit that part without affecting other parts. After all parts of your object are uploaded, Amazon S3 assembles these parts and creates the object. In general, when
your object size reaches 100 MB, you should consider using multipart uploads instead of uploading the object in a single operation.
Using multipart upload provides the following advantages:
Improved throughput—You can upload parts in parallel to improve throughput.
Quick recovery from any network issues—Smaller part size minimizes the impact of restarting a failed upload due to a network error.
Pause and resume object uploads—You can upload object parts over time. Once you initiate a multipart upload there is no expiry; you must explicitly complete or abort the multipart upload.
Begin an upload before you know the final object size—You can upload an object as you are creating it. Reference:
http://docs.aws.amazon.com/AmazonS3/latest/dev/uploadobjusingmpu.htmI

NEW QUESTION 226
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 230
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 Simp|eDB
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 235
You are setting up some IAM user policies and have also become aware that some services support resource-based permissions, which let you attach policies to the service's resources instead of to IAM users or groups. Which of the below statements is true in regards to resource-level permissions?

A. All services support resource-level permissions for all actions.
B. Resource-level permissions are supported by Amazon CloudFront
C. All services support resource-level permissions only for some actions.
D. Some services support resource-level permissions only for some action

Answer: D

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

The service is targeted at organizations with multiple users or systems that use AWS products such as Amazon EC2, Amazon RDS, and the AWS Management Console. With IAM, you can centrally manage users, security credentials such as access keys, and permissions that control which AWS resources users can access.
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, and Amazon SQS.
The resource-level permissions service supports IAM policies in which you can specify indMdual resources using Amazon Resource Names (ARNs) in the policy's Resource element.
Some services support resource-level permissions only for some actions.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_SpecificProducts.html

NEW QUESTION 240
In relation to AWS CIoudHSM, High-availability (HA) recovery is hands-off resumption by failed HA group members.
Prior to the introduction of this function, the HA feature provided redundancy and performance, but required that a failed/lost group member be reinstated.

A. automatically
B. periodically
C. manually
D. continuosly

Answer: C

Explanation:
In relation to AWS CIoudHS|VI, High-availability (HA) recovery is hands-off resumption by failed HA group members.
Prior to the introduction of this function, the HA feature provided redundancy and performance, but required that a failed/lost group member be manually reinstated.
Reference: http://docs.aws.amazon.com/cloudhsm/latest/userguide/ha-best-practices.html

NEW QUESTION 243
Having just set up your first Amazon Virtual Private Cloud (Amazon VPC) network, which defined a default network interface, you decide that you need to create and attach an additional network interface, known as an elastic network interface (ENI) to one of your instances. Which of the following statements is true regarding attaching network interfaces to your instances in your VPC?

A. You can attach 5 EN|s per instance type.
B. You can attach as many ENIs as you want.
C. The number of ENIs you can attach varies by instance type.
D. You can attach 100 ENIs total regardless of instance typ

Answer: C

Explanation:
Each instance in your VPC has a default network interface that is assigned a private IP address from the IP address range of your VPC. You can create and attach an additional network interface, known as an elastic network interface (ENI), to any instance in your VPC. The number of EN|s you can attach varies by instance type.

NEW QUESTION 246
You have set up an S3 bucket with a number of images in it and you have decided that you want anybody to be able to access these images, even anonymous users. To accomplish this you create a bucket policy. You will need to use an Amazon S3 bucket policy that specifies a in the principal element, which means anyone can access the bucket.

A. hash tag (#)
B. anonymous user
C. wildcard (*)
D. S3 user

Answer: C

Explanation:
You can use the AWS Policy Generator to create a bucket policy for your Amazon S3 bucket. You can then use the generated document to set your bucket policy by using the Amazon S3 console, by a number of third-party tools, or via your application.
You use an Amazon S3 bucket policy that specifies a wildcard (*) in the principal element, which means anyone can access the bucket. With anonymous access, anyone (including users without an AWS account) will be able to access the bucket.
Reference: http://docs.aws.amazon.com/IAM/|atest/UserGuide/iam-troubleshooting.htm|#d0e20565

NEW QUESTION 249
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 250
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 254
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 255
In Amazon Elastic Compute Cloud, which ofthe 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 259
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 263
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 266
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 "|nstance 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 270
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 273
Content and IV|edia Server is the latest requirement that you need to meet for a client.
The client has been very specific about his requirements such as low latency, high availability, durability, and access control. Potentially there will be millions of views on this server and because of "spiky" usage patterns, operations teams will need to provision static hardware, network, and management resources to support the maximum expected need. The Customer base will be initially low but is expected to grow and become more geographically distributed.
Which of the following would be a good solution for content distribution?

A. Amazon S3 as both the origin server and for caching
B. AWS Storage Gateway as the origin server and Amazon EC2 for caching
C. AWS CloudFront as both the origin server and for caching
D. Amazon S3 as the origin server and Amazon CloudFront for caching

Answer: D

Explanation:
As your customer base grows and becomes more geographically distributed, using a high- performance edge cache like Amazon CloudFront can provide substantial improvements in latency, fault tolerance, and cost.
By using Amazon S3 as the origin server for the Amazon CloudFront distribution, you gain the advantages of fast in-network data transfer rates, simple publishing/caching workflow, and a unified security framework.
Amazon S3 and Amazon CloudFront can be configured by a web service, the AWS Management Console, or a host of third-party management tools.
Reference:http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_media_02.pdf

NEW QUESTION 277
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 281
A user is trying to launch a similar EC2 instance from an existing instance with the option "Launch More like this". The AMI ofthe selected instance is deleted. What will happen in this case?

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

Answer: D

Explanation:
If the user has deregistered the AMI of an EC2 instance and is trying to launch a similar instance with the option "Launch more like this", AWS will throw an error saying that the AMI is deregistered or not available.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/launching-instance.html

NEW QUESTION 284
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/Iatest/PeeringGuide/vpc-peering-overview.htmI#vpc-peering-Ii mitations

NEW QUESTION 289
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 defaul

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 293
A friend wants you to set up a small BitTorrent storage area for him on Amazon S3. You tell him it is highly unlikely that AWS would allow such a thing in their infrastructure. However you decide to investigate. Which of the following statements best describes using BitTorrent with Amazon S3?

A. Amazon S3 does not support the BitTorrent protocol because it is used for pirated software.
B. You can use the BitTorrent protocol but only for objects that are less than 100 GB in size.
C. You can use the BitTorrent protocol but you need to ask AWS for specific permissions first.
D. You can use the BitTorrent protocol but only for objects that are less than 5 GB in siz

Answer: D

Explanation:
BitTorrent is an open, peer-to-peer protocol for distributing files. You can use the BitTorrent protocol to retrieve any publicly-accessible object in Amazon S3.

Amazon S3 supports the BitTorrent protocol so that developers can save costs when distributing content at high scale. Amazon S3 is useful for simple, reliable storage of any data. The default distribution mechanism for Amazon S3 data is via client/server download. In client/server distribution, the entire object is transferred point-to-point from Amazon S3 to every authorized user who requests that object. While client/server delivery is appropriate for a wide variety of use cases, it is not optimal for everybody. Specifically, the costs of client/server distribution increase linearly as the number of users downloading objects increases. This can make it expensive to distribute popular objects.

BitTorrent addresses this problem by recruiting the very clients that are downloading the object as distributors themselves: Each client downloads some pieces of the object from Amazon S3 and some from other clients, while simultaneously uploading pieces of the same object to other interested "peers." The benefit for publishers is that for large, popular files the amount of data actually supplied by Amazon S3 can be substantially lower than what it would have been sewing the same clients via client/server download. Less data transferred means lower costs for the publisher of the object.

Reference: http://docs.aws.amazon.com/AmazonS3/latest/dev/S3Torrent.html

NEW QUESTION 296
After a major security breach your manager has requested a report of all users and their credentials in AWS. You discover that in IAM you can generate and download a credential report that lists all users in your account and the status of their various credentials, including passwords, access keys, MFA devices, and signing certificates. Which following statement is incorrect in regards to the use of credential reports?

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

Answer: A

Explanation:
To access your AWS account resources, users must have credentials.
You can generate and download a credential report that lists all users in your account and the status of their various credentials, including passwords, access keys, MFA devices, and signing certificates. You can get a credential report using the AWS Management Console, the AWS CLI, or the IAM API.
You can use credential reports to assist in your auditing and compliance efforts. You can use the report to audit the effects of credential lifecycle requirements, such as password rotation. You can provide the report to an external auditor, or grant permissions to an auditor so that he or she can download the report directly.
You can generate a credential report as often as once every four hours. When you request a report, IAM first checks whether a report for the account has been generated within the past four hours. If so, the most recent report is downloaded. If the most recent report for the account is more than four hours old, or if there are no previous reports for the account, IAM generates and downloads a new report.
Credential reports are downloaded as comma-separated values (CSV) files.
You can open CSV files with common spreadsheet software to perform analysis, or you can build an application that consumes the CSV files programmatically and performs custom analysis. Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/credential-reports.html

NEW QUESTION 297
You have been given a scope to set up an AWS Media Sharing Framework for a new start up photo
sharing company similar to flickr. The first thing that comes to mind about this is that it will obviously need a huge amount of persistent data storage for this framework. Which of the following storage options would be appropriate for persistent storage?

A. Amazon Glacier or Amazon S3
B. Amazon Glacier or AWS Import/Export
C. AWS Import/Export or Amazon C|oudFront
D. Amazon EBS volumes or Amazon S3

Answer: D

Explanation:
Persistent storage-If you need persistent virtual disk storage similar to a physical disk drive for files or other data that must persist longer than the lifetime of a single Amazon EC2 instance, Amazon EBS volumes or Amazon S3 are more appropriate.
Reference: http://media.amazonwebservices.com/AWS_Storage_Options.pdf

NEW QUESTION 298
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 300
Which DNS name can only be resolved within Amazon EC2?

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

Answer: B

Explanation:
Only Internal DNS name can be resolved within Amazon EC2. Reference:

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

NEW QUESTION 301
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 u 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 302
You need to create a management network using network interfaces for a virtual private cloud (VPC) network. Which of the following statements is incorrect pertaining to Best Practices for Configuring Network Interfaces.

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

Answer: D

Explanation:
Best Practices for Configuring Network Interfaces
You can attach a network interface to an instance when it's running (hot attach), when it's stopped (warm attach), or when the instance is being launched (cold attach).
You can detach secondary (ethN) network interfaces when the instance is running or stopped. However, you can't detach the primary (eth0) interface.
You can attach a network interface in one subnet to an instance in another subnet in the same VPC, however, both the network interface and the instance must reside in the same Availability Zone.
When launching an instance from the CLI or API, you can specify the network interfaces to attach to the instance for both the primary (eth0) and additional network interfaces.
Launching an instance with multiple network interfaces automatically configures interfaces, private IP addresses, and route tables on the operating system of the instance.
A warm or hot attach of an additional network interface may require you to manually bring up the second interface, configure the private IP address, and modify the route table accordingly. (Instances running Amazon Linux automatically recognize the warm or hot attach and configure themselves.)
Attaching another network interface to an instance is not a method to increase or double the network bandwidth to or from the dual-homed instance.
Reference:
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html#use-network-and-security-applia nces-in-your-vpc

NEW QUESTION 305
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 sewers 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 sewers in the public subnet to read and write to MS SQL sewers in the private subnet. Port 1433 Inbound RDP traffic from the Microsoft Terminal Sewices gateway in the public private subnet. Port 3389 Reference:
http://docs.aws.amazon.com/AmazonVPC/Iatest/UserGuide/VPC_Appendix_NACLs.htm|#VPC_Appendi x_NAC Ls_Scenario_2

NEW QUESTION 310
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 314
In Amazon EC2, you are billed instance-hours when .

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

Answer: A

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

NEW QUESTION 318
Just when you thought you knew every possible storage option on AWS you hear someone mention Reduced Redundancy Storage (RRS) within Amazon S3.
What is the ideal scenario to use Reduced Redundancy Storage (RRS)?

A. Huge volumes of data
B. Sensitve data
C. Non-critical or reproducible data
D. Critical data

Answer: C

Explanation:
Reduced Redundancy Storage (RRS) is a new storage option within Amazon S3 that enables customers to reduce their costs by storing non-critical, reproducible data at lower levels of redundancy than Amazon S3's standard storage. RRS provides a lower cost, less durable, highly available storage option that is designed to sustain the loss of data in a single facility.
RRS is ideal for non-critical or reproducible data.
For example, RRS is a cost-effective solution for sharing media content that is durably stored elsewhere. RRS also makes sense if you are storing thumbnails and other resized images that can be easily reproduced from an original image.
Reference: https://aws.amazon.com/s3/faqs/

NEW QUESTION 323
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_To|erant_Applications.pdf

NEW QUESTION 325
You have some very sensitive data stored on AWS S3 and want to try every possible alternative to keeping it secure in regards to access control. What are the mechanisms available for access control on AWS S3?

A. (IAM) policies, Access Control Lists (ACLs), bucket policies, and query string authentication.
B. (IAM) policies, Access Control Lists (ACLs) and bucket policies.
C. Access Control Lists (ACLs), bucket policies, and query string authentication
D. (IAM) policies, Access Control Lists (ACLs), bucket policies, query string authentication and encryption.

Answer: A

Explanation:

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

NEW QUESTION 329
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 Nlobile Communication Service.
D. AWS Simple Queue Service.

Answer: A

Explanation:
Amazon Simple Notification Service (Amazon SNS) is a fast, filexible, 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 334
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.htmI

NEW QUESTION 335
......

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

    All transactions are protected by VeriSign!

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

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