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_|nstanceStraightToTerminated.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
Amazon EBS provides the ability to create backups of any Amazon EC2 volume into what is known as

A. snapshots
B. images
C. instance backups
D. mirrors

Answer: A

Explanation:
Amazon allows you to make backups of the data stored in your EBS volumes through snapshots that can later be used to create a new EBS volume.
Reference: http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/Storage.htmI

NEW QUESTION 5
A user is launching an EC2 instance in the US East region. Which of the below mentioned options is recommended by AWS with respect to the selection of the availability zone?

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

Answer: C

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

NEW QUESTION 6
A user is storing a large number of objects on AWS S3. The user wants to implement the search functionality among the objects. How can the user achieve this?

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

Answer: D

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

NEW QUESTION 7
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 8
True orfalsez A VPC contains multiple subnets, where each subnet can span multiple Availability Zones.

A. This is true only if requested during the set-up of VPC.
B. This is true.
C. This is false.
D. This is true only for US region

Answer: C

Explanation:
A VPC can span several Availability Zones. In contrast, a subnet must reside within a single Availability Zone.
Reference: https://aws.amazon.com/vpc/faqs/

NEW QUESTION 9
You are looking at ways to improve some existing infrastructure as it seems a lot of engineering resources are being taken up with basic management and monitoring tasks and the costs seem to be excessive.
You are thinking of deploying Amazon E|asticCache to help. Which of the following statements is true in regards to ElasticCache?

A. You can improve load and response times to user actions and queries however the cost associated with scaling web applications will be more.
B. You can't improve load and response times to user actions and queries but you can reduce the cost associated with scaling web applications.
C. You can improve load and response times to user actions and queries however the cost associated with scaling web applications will remain the same.
D. You can improve load and response times to user actions and queries and also reduce the cost associated with scaling web applications.

Answer: D

Explanation:

Amazon ElastiCache is a web service that makes it easy to deploy and run Memcached or Redis protocol-compliant server nodes in the cloud. Amazon ElastiCache improves the performance of web applications by allowing you to retrieve information from a fast, managed, in-memory caching system, instead of relying entirely on slower disk-based databases. The service simplifies and offloads the management, monitoring and operation of in-memory cache environments, enabling your engineering resources to focus on developing applications.
Using Amazon ElastiCache, you can not only improve load and response times to user actions and queries, but also reduce the cost associated with scaling web applications.
Reference: https://aws.amazon.com/elasticache/faqs/

NEW QUESTION 10
Do Amazon EBS volumes persist independently from the running life of an Amazon EC2 instance?

A. Yes, they do but only if they are detached from the instance.
B. No, you cannot attach EBS volumes to an instance.
C. No, they are dependent.
D. Yes, they d

Answer: D

Explanation:
An Amazon EBS volume behaves like a raw, unformatted, external block device that you can attach to a
single instance. The volume persists independently from the running life of an Amazon EC2 instance. Reference:
http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/Storage.html

NEW QUESTION 10
Your manager has just given you access to multiple VPN connections that someone else has recently set up between all your company's offices. She needs you to make sure that the communication between the VPNs is secure. Which of the following services would be best for providing a low-cost hub-and-spoke model for primary or backup connectMty between these remote offices?

A. Amazon C|oudFront
B. AWS Direct Connect
C. AWS C|oudHSM
D. AWS VPN CloudHub

Answer: D

Explanation:
If you have multiple VPN connections, you can provide secure communication between sites using the
AWS VPN CloudHub. The VPN CloudHub operates on a simple hub-and-spoke model that you can use with or without a VPC. This design is suitable for customers with multiple branch offices and existing Internet connections who would like to implement a convenient, potentially low-cost hub-and-spoke model for primary or backup connectMty between these remote offices.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPN_CloudHub.html

NEW QUESTION 12
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 16
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 17
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 19
You are in the process of creating a Route 53 DNS failover to direct traffic to two EC2 zones. Obviously, if one fails, you would like Route 53 to direct traffic to the other region. Each region has an ELB with some instances being distributed. What is the best way for you to configure the Route 53 health check?

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

Answer: D

Explanation:
With DNS Failover, Amazon Route 53 can help detect an outage of your website and redirect your end users to alternate locations where your application is operating properly. When you enable this feature, Route 53 uses health checks-regularly making Internet requests to your application's endpoints from multiple locations around the world-to determine whether each endpoint of your application is up or down.
To enable DNS Failover for an ELB endpoint, create an Alias record pointing to the ELB and set the "Evaluate Target Health" parameter to true. Route 53 creates and manages the health checks for your ELB automatically. You do not need to create your own Route 53 health check of the ELB. You also do not need to associate your resource record set for the ELB with your own health check, because Route 53 automatically associates it with the health checks that Route 53 manages on your behalf. The ELB health check will also inherit the health of your backend instances behind that ELB.
Reference:
http://aws.amazon.com/about-aws/whats-new/2013/05/30/amazon-route-53-adds-elb-integration-for-dns- fai|over/

NEW QUESTION 21
A user wants to use an EBS-backed Amazon EC2 instance for a temporary job. Based on the input data, the job is most likely to finish within a week. Which of the following steps should be followed to terminate the instance automatically once the job is finished?

A. Configure the EC2 instance with a stop instance to terminate it.
B. Configure the EC2 instance with ELB to terminate the instance when it remains idle.
C. Configure the CIoudWatch alarm on the instance that should perform the termination action once the instance is idle.
D. Configure the Auto Scaling schedule actMty that terminates the instance after 7 day

Answer: C

Explanation:
Auto Scaling can start and stop the instance at a pre-defined time. Here, the total running time is unknown. Thus, the user has to use the CIoudWatch alarm, which monitors the CPU utilization. The user can create an alarm that is triggered when the average CPU utilization percentage has been lower than 10 percent for 24 hours, signaling that it is idle and no longer in use. When the utilization is below the threshold limit, it will terminate the instance as a part of the instance action.
Reference: http://docs.aws.amazon.com/AmazonCIoudWatch/|atest/Deve|operGuide/UsingAIarmActions.html

NEW QUESTION 23
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 25
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 29
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 32
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 34
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 36
You are checking the workload on some of your General Purpose (SSD) and Provisioned IOPS (SSD) volumes and it seems that the I/O latency is higher than you require. You should probably check the to make sure that your application is not trying to drive more IOPS than you have provisioned.

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

Answer: C

Explanation:
In EBS workload demand plays an important role in getting the most out of your General Purpose (SSD) and Provisioned IOPS (SSD) volumes. In order for your volumes to deliver the amount of IOPS that are available, they need to have enough I/O requests sent to them. There is a relationship between the demand on the volumes, the amount of IOPS that are available to them, and the latency of the request (the amount of time it takes for the I/O operation to complete).
Latency is the true end-to-end client time of an I/O operation; in other words, when the client sends a IO, how long does it take to get an acknowledgement from

the storage subsystem that the IO read or write is complete.
If your I/O latency is higher than you require, check your average queue length to make sure that your application is not trying to drive more IOPS than you have provisioned. You can maintain high IOPS while keeping latency down by maintaining a low average queue length (which is achieved by provisioning more IOPS for your volume).
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-workload-demand.htmI

NEW QUESTION 39
Which of the below mentioned options is not available when an instance is launched by Auto Scaling with EC2 Classic?

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

Answer: B

Explanation:
Auto Scaling supports both EC2 classic and EC2-VPC. When an instance is launched as a part of EC2 classic, it will have the public IP and DNS as well as the private IP and DNS.
Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/GettingStartedTutorial.html

NEW QUESTION 44
In DynamoDB, could you use IAM to grant access to Amazon DynamoDB resources and API actions?

A. In DynamoDB there is no need to grant access
B. Depended to the type of access
C. No
D. Yes

Answer: D

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

NEW QUESTION 48
Much of your company's data does not need to be accessed often, and can take several hours for retrieval time, so it's stored on Amazon Glacier. However someone within your organization has expressed concerns that his data is more sensitive than the other data, and is wondering whether the high level of encryption that he knows is on S3 is also used on the much cheaper Glacier service. Which of the following statements would be most applicable in regards to this concern?

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

Answer: C

Explanation:
Like Amazon S3, the Amazon Glacier service provides low-cost, secure, and durable storage. But where S3 is designed for rapid retrieval, Glacier is meant to be used as an archival service for data that is not accessed often, and for which retrieval times of several hours are suitable.
Amazon Glacier automatically encrypts the data using AES-256 and stores it durably in an immutable form. Amazon Glacier is designed to provide average annual durability of 99.999999999% for an archive. It stores each archive in multiple facilities and multiple devices. Unlike traditional systems which can require laborious data verification and manual repair, Glacier performs regular, systematic data integrity checks, and is built to be automatically self-healing.
Reference: http://d0.awsstatic.com/whitepapers/Security/AWS%20Security%20Whitepaper.pdf

NEW QUESTION 51
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 52
In Amazon EC2 Container Service, are other container types supported?

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

Answer: C

Explanation:
In Amazon EC2 Container Service, Docker is the only container platform supported by EC2 Container Service presently.
Reference: http://aws.amazon.com/ecs/faqs/

NEW QUESTION 56
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 60
You need to set up a complex network infrastructure for your organization that will be reasonably easy to deploy, replicate, control, and track changes on. Which AWS service would be best to use to help you accomplish this?

A. AWS Import/Export
B. AWS CIoudFormation
C. Amazon Route 53
D. Amazon CIoudWatch

Answer: B

Explanation:
AWS CIoudFormation is a service that helps you model and set up your Amazon Web Services resources so that you can spend less time managing those resources and more time focusing on your applications that run in AWS. You create a template that describes all the AWS resources that you want (like Amazon EC2 instances or Amazon RDS DB instances), and AWS CIoudFormation takes care of provisioning and configuring those resources for you. You don't need to indMdually create and configure AWS resources
and figure out what's dependent on what. AWS CIoudFormation handles all of that.
Reference: http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/WeIcome.htmI

NEW QUESTION 63
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 67
In an experiment, if the minimum size for an Auto Scaling group is 1 instance, which of the following statements holds true when you terminate the running instance?

A. Auto Scaling must launch a new instance to replace it.
B. Auto Scaling will raise an alarm and send a notification to the user for action.

C. Auto Scaling must configure the schedule actMty that terminates the instance after 5 days.
D. Auto Scaling will terminate the experimen

Answer: A

Explanation:
If the minimum size for an Auto Scaling group is 1 instance, when you terminate the running instance, Auto Scaling must launch a new instance to replace it.
Reference:http://docs.aws.amazon.com/AutoScaling/latest/Deve|operGuide/AS_Concepts.htmI

NEW QUESTION 70
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 72
You are running PostgreSQL on Amazon RDS and it seems to be all running smoothly deployed in one availability zone. A database administrator asks you if DB instances running PostgreSQL support MuIti-AZ deployments. What would be a correct response to this QUESTION ?

A. Yes.
B. Yes but only for small db instances.
C. No.
D. Yes but you need to request the service from AW

Answer: A

Explanation:
Amazon RDS supports DB instances running several versions of PostgreSQL. Currently we support PostgreSQL versions 9.3.1, 9.3.2, and 9.3.3. You can create DB instances and DB snapshots,
point-in-time restores and backups.
DB instances running PostgreSQL support MuIti-AZ deployments, Provisioned IOPS, and can be created inside a VPC. You can also use SSL to connect to a DB instance running PostgreSQL.
You can use any standard SQL client application to run commands for the instance from your client computer. Such applications include pgAdmin, a popular Open Source administration and development tool for PostgreSQL, or psql, a command line utility that is part of a PostgreSQL installation. In order to deliver a managed service experience, Amazon RDS does not provide host access to DB instances, and it restricts access to certain system procedures and tables that require advanced prMleges. Amazon RDS supports access to databases on a DB instance using any standard SQL client application. Amazon RDS does not allow direct host access to a DB instance via Telnet or Secure Shell (SSH).
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_PostgreSQL.htmI

NEW QUESTION 74
Which of the following AWS CLI commands is syntactically incorrect?
1. $ aws ec2 describe-instances
2. $ aws ec2 start-instances --instance-ids i-1348636c
3. $ aws sns publish --topic-arn arn:aws:sns:us-east-1:546419318123:OperationsError -message "Script Failure"
4. $ aws sqs receive-message --queue-urI https://queue.amazonaws.com/546419318123/Test

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

Answer: A

Explanation:
The following CLI command is missing a hyphen before "-message".
aws sns publish --topic-arn arn:aws:sns:us-east-1:546419318123:OperationsError -message "Script Failure"
It has been added below in red
aws sns publish --topic-arn arn:aws:sns:us-east-1:546419318123:OperationsError ---message "Script Failure"
Reference: http://aws.amazon.com/c|i/

NEW QUESTION 76
You are architecting an auto-scalable batch processing system using video processing pipelines and Amazon Simple Queue Service (Amazon SQS) for a customer. You are unsure of the limitations of SQS and need to find out. What do you think is a correct statement about the limitations of Amazon SQS?

A. It supports an unlimited number of queues but a limited number of messages per queue for each user but automatically deletes messages that have been in the queue for more than 4 weeks.
B. It supports an unlimited number of queues and unlimited number of messages per queue for each user but automatically deletes messages that have been in the queue for more than 4 days.
C. It supports an unlimited number of queues but a limited number of messages per queue for each user but automatically deletes messages that have been in the queue for more than 4 days.
D. It supports an unlimited number of queues and unlimited number of messages per queue for each user but automatically deletes messages that have been in

the queue for more than 4 weeks.

Answer: B

Explanation:
Amazon Simple Queue Service (Amazon SQS) is a messaging queue service that handles message or workflows between other components in a system. Amazon SQS supports an unlimited number of queues and unlimited number of messages per queue for each user. Please be aware that Amazon SQS automatically deletes messages that have been in the queue for more than 4 days.
Reference: http://aws.amazon.com/documentation/sqs/

NEW QUESTION 81
An online gaming site asked you if you can deploy a database that is a fast, highly scalable NoSQL database service in AWS for a new site that he wants to build. Which database should you recommend?

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

Answer: A

Explanation:
Amazon DynamoDB is ideal for database applications that require very low latency and predictable performance at any scale but don't need complex querying capabilities like joins or transactions. Amazon DynamoDB is a fully-managed NoSQL database service that offers high performance, predictable throughput and low cost. It is easy to set up, operate, and scale.
With Amazon DynamoDB, you can start small, specify the throughput and storage you need, and easily scale your capacity requirements on the fly. Amazon DynamoDB automatically partitions data over a number of servers to meet your request capacity. In addition, DynamoDB automatically replicates your data synchronously across multiple Availability Zones within an AWS Region to ensure high-availability and data durability.
Reference: https://aws.amazon.com/running_databases/#dynamodb_anchor

NEW QUESTION 86
You log in to IAM on your AWS console and notice the following message. "Delete your root access keys." Why do you think IAM is requesting this?

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

Answer: D

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

NEW QUESTION 89
A customer enquires about whether all his data is secure on AWS and is especially concerned about Elastic Map Reduce (EMR) so you need to inform him of some of the security features in place for AWS. Which of the below statements would be an incorrect response to your customers enquiry?

A. Amazon ENIR customers can choose to send data to Amazon S3 using the HTTPS protocol for secure transmission.
B. Amazon S3 provides authentication mechanisms to ensure that stored data is secured against unauthorized access.
C. Every packet sent in the AWS network uses Internet Protocol Security (IPsec).
D. Customers may encrypt the input data before they upload it to Amazon S3.

Answer: C

Explanation:
Amazon S3 provides authentication mechanisms to ensure that stored data is secured against unauthorized access. Unless the customer who is uploading the data specifies otherwise, only that customer can access the data. Amazon EMR customers can also choose to send data to Amazon S3 using the HTTPS protocol for secure transmission. In addition, Amazon EMR always uses HTTPS to send data between Amazon S3 and Amazon EC2. For added security, customers may encrypt the input data before they upload it to Amazon S3 (using any common data compression tool); they then need to add a decryption step to the beginning of their cluster when Amazon EMR fetches the data from Amazon S3. Reference: https://aws.amazon.com/elasticmapreduce/faqs/

NEW QUESTION 91
You are in the process of building an online gaming site for a client and one of the requirements is that it must be able to process vast amounts of data easily. Which AWS Service would be very helpful in processing all this data?

A. Amazon S3
B. AWS Data Pipeline
C. AWS Direct Connect
D. Amazon EMR

Answer: D

Explanation:

Managing and analyzing high data volumes produced by online games platforms can be difficult. The back-end infrastructures of online games can be challenging to maintain and operate. Peak usage periods, multiple players, and high volumes of write operations are some of the most common problems that operations teams face.

Amazon Elastic MapReduce (Amazon EMR) is a service that processes vast amounts of data easily. Input data can be retrieved from web server logs stored on Amazon S3 or from player data stored in Amazon DynamoDB tables to run analytics on player behavior, usage patterns, etc. Those results can be stored again on Amazon S3, or inserted in a relational database for further analysis with classic business intelligence tools.

Reference: http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_games_10.pdf

NEW QUESTION 93

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

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

A. For security reasons.
B. Hardware restrictions.
C. Public (IPV4) internet addresses are a scarce resource.
D. There are only 5 network interfaces per instanc

Answer: C

Explanation:

Public (IPV4) internet addresses are a scarce resource. There is only a limited amount of public IP space available, and Amazon EC2 is committed to helping use that space efficiently.

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

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

NEW QUESTION 94

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

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

Answer: D

Explanation:

Amazon RDS provides high availability and failover support for DB instances using Multi-AZ deployments. Multi-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 97

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 101

Having set up a website to automatically be redirected to a backup website if it fails, you realize that there are different types of failovers that are possible. You need all your resources to be available the majority of the time. Using Amazon Route 53 which configuration would best suit this requirement?

A. Active-active failover.
B. Non
C. Route 53 can't failover.
D. Active-passive failover.
E. Active-active-passive and other mixed configuration

Answer: A

Explanation:

You can set up a variety of failover configurations using Amazon Route 53 alias: weighted, latency, geolocation routing, and failover resource record sets.

Active-active failover: Use this failover configuration when you want all of your resources to be available the majority of the time. When a resource becomes unavailable, Amazon Route 53 can detect that it's unhealthy and stop including it when responding to queries.

Active-passive failover: Use this failover configuration when you want a primary group of resources to be available the majority of the time and you want a secondary group of resources to be on standby in case all of the primary resources become unavailable. When responding to queries, Amazon Route 53 includes only the healthy primary resources. If all of the primary resources are unhealthy, Amazon Route 53 begins to include only the healthy secondary resources in response to DNS queries.

Active-active-passive and other mixed configurations: You can combine alias and non-alias resource record sets to produce a variety of Amazon Route 53

behaviors.
Reference: http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover.html

NEW QUESTION 105
AWS CloudFormation is a service that helps you model and set up your Amazon Web Services resources so that you can spend less time managing those resources and more time focusing on your applications that run in AWS. You create a template that describes all the AWS resources that you want (like Amazon EC2 instances or Amazon RDS DB instances), and AWS CloudFormation takes care of provisioning and configuring those resources for you. What formatting is required for this template?

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

Answer: A

Explanation:
You can write an AWS CloudFormation template (a JSON-formatted document) in a text editor or pick an existing template. The template describes the resources you want and their settings. For example,
suppose you want to create an Amazon EC2. Your template can declare an instance Amazon EC2 and describe its properties, as shown in the following example:
{
"AWSTemp|ateFormatVersion" : "2010-09-O9",
"Description" : "A simple Amazon EC2 instance", "Resources" : {
"MyEC2Instance" : {
"Type" : "AWS::EC2::Instance", "Properties" : {
"Image|d" : "ami-2f726546", "|nstanceType" : "t1.micro"
}
}
}
}
Reference:
http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/cfn-whatis-howdoesitwork.html

NEW QUESTION 106
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 107
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 111
You have been storing massive amounts of data on Amazon Glacier for the past 2 years and now start to wonder if there are any limitations on this. What is the correct answer to your QUESTION ?

A. The total volume of data is limited but the number of archives you can store are unlimited.
B. The total volume of data is unlimited but the number of archives you can store are limited.
C. The total volume of data and number of archives you can store are unlimited.
D. The total volume of data is limited and the number of archives you can store are limite

Answer: C

Explanation:
An archive is a durably stored block of information. You store your data in Amazon Glacier as archives. You may upload a single file as an archive, but your costs will be lower if you aggregate your data. TAR and ZIP are common formats that customers use to aggregate multiple files into a single file before uploading to Amazon Glacier.
The total volume of data and number of archives you can store are unlimited. IndMdual Amazon Glacier archives can range in size from 1 byte to 40 terabytes. The largest archive that can be uploaded in a single upload request is 4 gigabytes.
For items larger than 100 megabytes, customers should consider using the Multipart upload capability. Archives stored in Amazon Glacier are immutable, i.e. archives can be uploaded and deleted but cannot be edited or overwritten.
Reference: https://aws.amazon.com/gIacier/faqs/

NEW QUESTION 112
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_CIoud_Best_Practices.pdf

NEW QUESTION 115
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 118
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 121
You are building a system to distribute confidential documents to employees. Using CIoudFront, what method could be used to serve content that is stored in S3, but not publically accessible from S3 directly?

A. Add the CIoudFront account security group "amazon-cf/amazon-cf-sg" to the appropriate S3 bucket policy.
B. Create a S3 bucket policy that lists the C|oudFront distribution ID as the Principal and the target bucket as the Amazon Resource Name (ARN).
C. Create an Identity and Access Management (IAM) User for CloudFront and grant access to the objects in your S3 bucket to that IAM User.
D. Create an Origin Access Identity (OAI) for CIoudFront and grant access to the objects in your S3 bucket to that OAI.

Answer: D

Explanation:
You restrict access to Amazon S3 content by creating an origin access identity, which is a special CloudFront user. You change Amazon S3 permissions to give the origin access identity permission to access your objects, and to remove permissions from everyone else. When your users access your Amazon S3 objects using CloudFront URLs, the CloudFront origin access identity gets the objects on your users' behalf. If your users try to access objects using Amazon S3 URLs, they're denied access. The origin access identity has permission to access objects in your Amazon S3 bucket, but users don't. Reference: http://docs.aws.amazon.com/AmazonCloudFront/latest/Deve|operGuide/private-content-restricting-acces s-to-s3.html

NEW QUESTION 123
A user has created photo editing software and hosted it on EC2. The software accepts requests from the user about the photo format and resolution and sends a message to S3 to enhance the picture accordingly. Which of the below mentioned AWS services will help make a scalable software with the AWS infrastructure in this scenario?

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

Answer: B

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

NEW QUESTION 128
An accountant asks you to design a small VPC network for him and, due to the nature of his business, just needs something where the workload on the network will be low, and dynamic data will be accessed infrequently. Being an accountant, low cost is also a major factor. Which EBS volume type would best suit his requirements?

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

Answer: A

Explanation:
You can choose between three EBS volume types to best meet the needs of their workloads: General Purpose (SSD), Provisioned IOPS (SSD), and Magnetic. General Purpose (SSD) is the new, SSD-backed, general purpose EBS volume type that we recommend as the default choice for customers. General Purpose (SSD) volumes are suitable for a broad range of workloads, including small to medium sized databases, development and test environments, and boot volumes. Provisioned IOPS (SSD) volumes offer storage with consistent and low-latency performance, and are designed for I/O intensive applications such as large relational or NoSQL databases. Magnetic volumes provide the lowest cost per gigabyte of all EBS volume types. Magnetic volumes are ideal for workloads where data is accessed infrequently, and applications where the lowest storage cost is important.
Reference: https://aws.amazon.com/ec2/faqs/

NEW QUESTION 133
A user has launched one EC2 instance in the US East region and one in the US West region. The user has launched an RDS instance in the US East region. How can the user configure access from both the EC2 instances to RDS?

A. It is not possible to access RDS of the US East region from the US West region
B. Configure the US West region's security group to allow a request from the US East region's instance and configure the RDS security group's ingress rule for the US East EC2 group
C. Configure the security group of the US East region to allow traffic from the US West region's instance and configure the RDS security group's ingress rule for the US East EC2 group
D. Configure the security group of both instances in the ingress rule of the RDS security group

Answer: C

Explanation:
The user cannot authorize an Amazon EC2 security group if it is in a different AWS Region than the RDS
DB instance. The user can authorize an IP range or specify an Amazon EC2 security group in the same region that refers to an IP address in another region. In this case allow IP of US West inside US East's security group and open the RDS security group for US East region.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithSecurityGroups.html

NEW QUESTION 136
In Amazon EC2, if your EBS volume stays in the detaching state, you can force the detachment by clicking .

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

Answer: A

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

NEW QUESTION 141
An organization has a statutory requirement to protect the data at rest for data stored in EBS volumes. Which of the below mentioned options can the organization use to achieve data protection?

A. Data replication.
B. Data encryption.
C. Data snapshot.
D. All the options listed her

Answer: D

Explanation:
For protecting the Amazon EBS data at REST, the user can use options, such as Data Encryption (Windows / Linux / third party based), Data Replication (AWS internally replicates data for redundancy),
and Data Snapshot (for point in time backup).
Reference: http://media.amazonwebservices.com/AWS_Security_Best_Practices.pdf

NEW QUESTION 146
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 149
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 151
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 154
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 156
Which one of the below is not an AWS Storage Service?

A. Amazon S3
B. Amazon Glacier
C. Amazon CIoudFront
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 157
A user has deployed an application on his private cloud. The user is using his own monitoring tool. He wants to configure it so that whenever there is an error, the monitoring tool will notify him via SMS. Which of the below mentioned AWS services will help in this scenario?

A. AWS SES
B. AWS SNS
C. None because the user infrastructure is in the private cloud.
D. AWS SMS

Answer: B

Explanation:
Amazon Simple Notification Service (Amazon SNS) is a fast, filexible, and fully managed push messaging service. Amazon SNS can be used to make push notifications to mobile devices. Amazon SNS can
deliver notifications by SMS text message or email to the Amazon Simple Queue Service (SQS) queues or to any HTTP endpoint. In this case user can use the SNS apis to send SMS.
Reference: http://aws.amazon.com/sns/

NEW QUESTION 161
After setting up an EC2 security group with a cluster of 20 EC2 instances, you find an error in the security group settings. You quickly make changes to the security group settings. When will the changes to the settings be effective?

A. The settings will be effective immediately for all the instances in the security group.
B. The settings will be effective only when all the instances are restarted.
C. The settings will be effective for all the instances only after 30 minutes.
D. The settings will be effective only for the new instances added to the security grou

Answer: A

Explanation:
Amazon Redshift applies changes to a cluster security group immediately. So if you have associated the cluster security group with a cluster, inbound cluster access rules in the updated cluster security group apply immediately.
Reference: http://docs.aws.amazon.com/redshift/latest/mgmt/working-with-security-groups.htm|

NEW QUESTION 163
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/DeveIoperGuide/Tutorials.html

NEW QUESTION 167
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 169
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 c4.8xlarge instance is 10 Gigabit. Reference: http://aws.amazon.com/ec2/instance-types/

NEW QUESTION 171
Doug has created a VPC with CIDR 10.201.0.0/16 in his AWS account. In this VPC he has created a public subnet with CIDR block 10.201.31.0/24. While launching a new EC2 from the console, he is not able to assign the private IP address 10.201.31.6 to this instance. Which is the most likely reason for this issue?

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

Answer: B

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

NEW QUESTION 175
You need to create a JSON-formatted text file for AWS CIoudFormation. This is your first template and the only thing you know is that the templates include several major sections but there is only one that is required for it to work. What is the only section required?

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

Answer: C

Explanation:
AWS CIoudFormation is a service that helps you model and set up your Amazon Web Services resources so that you can spend less time managing those resources and more time focusing on your applications that run in AWS. You create a template that describes all the AWS resources that you want (like Amazon EC2 instances or Amazon RDS DB instances), and AWS CIoudFormation takes care of provisioning and configuring those resources for you.
A template is a JSON-formatted text file that describes your AWS infrastructure. Templates include several major sections.
The Resources section is the only section that is required.
The first character in the template must be an open brace ({), and the last character must be a closed brace (}). The following template fragment shows the template structure and sections.
Reference: http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/template-anatomy.html

NEW QUESTION 180
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 181
A user is hosting a website in the US West-1 region. The website has the highest client base from the Asia-Pacific (Singapore / Japan) region. The application is accessing data from S3 before serving it to client. Which of the below mentioned regions gives a better performance for S3 objects?

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

Answer: D

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

NEW QUESTION 185
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 189
You have three Amazon EC2 instances with Elastic IP addresses in the US East (Virginia) region, and you want to distribute requests across all three IPs evenly for users for whom US East (Virginia) is the appropriate region.
How many EC2 instances would be sufficient to distribute requests in other regions?

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

Answer: D

Explanation:
If your application is running on Amazon EC2 instances in two or more Amazon EC2 regions, and if you have more than one Amazon EC2 instance in one or more regions, you can use latency-based routing to route traffic to the correct region and then use weighted resource record sets to route traffic to instances within the region based on weights that you specify.
For example, suppose you have three Amazon EC2 instances with Elastic IP addresses in the US East (Virginia) region and you want to distribute requests across all three IPs evenly for users for whom US East (Virginia) is the appropriate region. Just one Amazon EC2 instance is sufficient in the other regions, although you can apply the same technique to many regions at once.
Reference: http://docs.aws.amazon.com/Route53/latest/DeveIoperGuide/Tutorials.html

NEW QUESTION 190
A major client who has been spending a lot of money on his internet service provider asks you to set up an AWS Direct Connection to try and save him some money. You know he needs high-speed connectMty. Which connection port speeds are available on AWS Direct Connect?

A. 500Mbps and 1Gbps
B. 1Gbps and 10Gbps
C. 100Mbps and 1Gbps
D. 1Gbps

Answer: B

Explanation:
AWS Direct Connect is a network service that provides an alternative to using the internet to utilize AWS cloud services.
Using AWS Direct Connect, data that would have previously been transported over the Internet can now be delivered through a private network connection between AWS and your datacenter or corporate network.
1Gbps and 10Gbps ports are available. Speeds of 50Mbps, 100Mbps, 200Mbps, 300Mbps, 400Mbps, and 500Mbps can be ordered from any APN partners supporting AWS Direct Connect.
Reference: https://aws.amazon.com/directconnect/faqs/

NEW QUESTION 192
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 196
The AWS CIoudHSM 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 CIoudHSNI 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.htmI

NEW QUESTION 198
You have just set up yourfirst Elastic Load Balancer (ELB) but it does not seem to be configured properly. You discover that before you start using ELB, you have to configure the listeners for your load balancer. Which protocols does ELB use to support the load balancing of applications?

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

Answer: C

Explanation:
Before you start using Elastic Load BaIancing(ELB), you have to configure the listeners for your load balancer. A listener is a process that listens for connection requests. It is configured with a protocol and a port number for front-end (client to load balancer) and back-end (load balancer to back-end instance) connections. Elastic Load Balancing supports the load balancing of applications using HTTP, HTTPS (secure HTTP), TCP, and SSL (secure TCP) protocols. The HTTPS uses the SSL protocol to establish secure connections over the HTTP layer. You can also use SSL protocol to establish secure connections over the TCP layer.
The acceptable ports for both HTTPS/SSL and HTTP/TCP connections are 25, 80, 443, 465, 587, and
1024-65535.
Reference:
http://docs.aws.amazon.com/E|asticLoadBaIancing/latest/DeveIoperGuide/elb-listener-config.htmI

NEW QUESTION 199
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 CIoudWatch 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 CIoudWatch
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 CIoudWatch is a monitoring service for AWS cloud resources and the applications you run on AWS. You can use Amazon CIoudWatch to collect and track metrics, collect and monitor log files, and set
alarms.
When you create an alarm, you first choose the Amazon CIoudWatch 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/cIoudwatch/faqs/

NEW QUESTION 201
When does the billing of an Amazon EC2 system begin?

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

Answer: D

Explanation:
Billing commences when Amazon EC2 initiates the boot sequence of an AM instance. Billing ends when the instance terminates, which could occur through a web services command, by running "shutdown -h", or through instance failure. When you stop an instance, Amazon shuts it down but doesn/Et charge hourly usage for

a stopped instance, or data transfer fees, but charges for the storage for any Amazon EBS volumes.
Reference: http://aws.amazon.com/ec2/faqs/

NEW QUESTION 206
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.html

NEW QUESTION 209
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 212
Mike is appointed as Cloud Consultant in Netcrak Inc. Netcrak has the following VPCs set-up in the US East Region:
A VPC with CIDR block 10.10.0.0/16, a subnet in that VPC with CIDR block 10.10.1.0/24 A VPC with CIDR block 10.40.0.0/16, a subnet in that VPC with CIDR block 10.40.1.0/24
Netcrak Inc is trying to establish network connection between two subnets, a subnet with CIDR block 10.10.1.0/24 and another subnet with CIDR block 10.40.1.0/24. Which one of the following solutions should Mke recommend to Netcrak Inc?

A. Create 2 Virtual Private Gateways and configure one with each VPC.
B. Create one EC2 instance in each subnet, assign Elastic IPs to both instances, and configure a set up Site-to-Site VPN connection between both EC2 instances.
C. Create a VPC Peering connection between both VPCs.
D. Create 2 Internet Gateways, and attach one to each VP

Answer: C

Explanation:
A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IP addresses. EC2 instances in either VPC can communicate with each other as if they are within the same network. You can create a VPC peering connection between your own VPCs, or with a VPC in another AWS account within a single region.
AWS uses the existing infrastructure of a VPC to create a VPC peering connection; it is neither a gateway nor a VPN connection, and does not rely on a separate piece of physical hardware.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpc-peering.htm|

NEW QUESTION 216
You're trying to delete an SSL certificate from the IAM certificate store, and you're getting the message "Certificate: <certificate-id> is being used by CIoudFront." Which of the following statements is probably the reason why you are getting this error?

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

Answer: A

Explanation:
CIoudFront is a web service that speeds up distribution of your static and dynamic web content, for example, .htmI, .css, .php, and image files, to end users.
Every CIoudFront web distribution must be associated either with the default CloudFront certificate or with a custom SSL certificate. Before you can delete an SSL

certificate, you need to either rotate SSL certificates (replace the current custom SSL certificate with another custom SSL certificate) or revert from using a custom SSL certificate to using the default CloudFront certificate.
Reference: http://docs.aws.amazon.com/AmazonCIoudFront/latest/Deve|operGuide/Troubleshooting.htm|

NEW QUESTION 220
You need to set up security for your VPC and you know that Amazon VPC provides two features that you can use to increase security for your VPC: Security groups and network access control lists (ACLs). You start to look into security groups first. Which statement below is incorrect in relation to security groups?

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

Answer: C

Explanation:
Amazon VPC provides two features that you can use to increase security for your VPC:
Security groups—Act as a firewall for associated Amazon EC2 instances, controlling both inbound and outbound traffic at the instance level and supports allow rules only.
Network access control lists (ACLs)—Act as a firewall for associated subnets, controlling both inbound and outbound traffic at the subnet level and supports allow rules and deny rules.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Security.html

NEW QUESTION 224
A user wants to increase the durability and availability of the EBS volume. Which of the below mentioned actions should he perform?

A. Take regular snapshots.
B. Create an AMI.
C. Create EBS with higher capacity.
D. Access EBS regularl

Answer: A

Explanation:
In Amazon Web Services, Amazon EBS volumes that operate with 20 GB or less of modified data since their most recent snapshot can expect an annual failure rate (AFR) between 0.1% and 0.5%. For this reason, to maximize both durability and availability of their Amazon EBS data, the user should frequently create snapshots of the Amazon EBS volumes.
Reference: http://media.amazonwebservices.com/AWS_Storage_Options.pdf

NEW QUESTION 226
A for a VPC is a collection of subnets (typically private) that you may want to designate for your backend RDS DB Instances.

A. DB Subnet Set
B. RDS Subnet Group
C. DB Subnet Group
D. DB Subnet Collection

Answer: C

Explanation:
DB Subnet Groups are a set of subnets (one per Availability Zone of a particular region) designed for your DB instances that reside in a VPC. They make easy to manage Multi-AZ deployments as well as the conversion from a Single-AZ to a Mut|i-AZ one.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.RDSVPC.htmI

NEW QUESTION 228
Amazon Elastic Load Balancing is used to manage traffic on a fileet of Amazon EC2 instances, distributing traffic to instances across all availability zones within a region. Elastic Load Balancing has all the advantages of an on-premises load balancer, plus several security benefits.
Which of the following is not an advantage of ELB over an on-premise load balancer?

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

Answer: A

Explanation:
Amazon Elastic Load Balancing is used to manage traffic on a fileet of Amazon EC2 instances, distributing traffic to instances across all availability zones within a region. Elastic Load Balancing has all the advantages of an on-premises load balancer, plus several security benefits:
Takes over the encryption and decryption work from the Amazon EC2 instances and manages it centrally on the load balancer
Offers clients a single point of contact, and can also serve as the first line of defense against attacks on your network
When used in an Amazon VPC, supports creation and management of security groups associated with your Elastic Load Balancing to provide additional networking and security options
Supports end-to-end traffic encryption using TLS (previously SSL) on those networks that use secure HTTP (HTTPS) connections. When TLS is used, the TLS server certificate used to terminate client connections can be managed centrally on the load balancer, rather than on every indMdual instance. Reference:
http://d0.awsstatic.com/whitepapers/Security/AWS%20Security%20Whitepaper.pdf

NEW QUESTION 229

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 233
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 236
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 239
A user has created an ELB with the availability zone US-East-1A. The user wants to add more zones to ELB to achieve High Availability. How can the user add more zones to the existing ELB?

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

Answer: D

Explanation:
The user has created an Elastic Load Balancer with the availability zone and wants to add more zones to the existing ELB. The user can do so in two ways:
From the console or CLI, add new zones to ELB;
Launch instances in a separate AZ and add instances to the existing ELB. Reference:
http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/enable-disable-az.html

NEW QUESTION 244
In AWS CloudHSM, in addition to the AWS recommendation that you use two or more HSM appliances in a high-availability configuration to prevent the loss of keys and data, you can also perform a remote backup/restore of a Luna SA partition if you have purchased a:

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

Answer: B

Explanation:
In AWS CloudHSM, you can perform a remote backup/restore of a Luna SA partition if you have purchased a Luna Backup HSM.
Reference: http://docs.aws.amazon.com/cloudhsm/latest/userguide/cloud-hsm-backup-restore.html

NEW QUESTION 245
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 246
An EC2 instance is connected to an ENI (Elastic Network Interface) in one subnet. What happens when you attach an ENI of a different subnet to this EC2 instance?

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

Answer: B

Explanation:
AWS allows you create an elastic network interface (ENI), attach an ENI to an EC2 instance, detach an ENI from an EC2 instance and attach this ENI to another EC2 instance. The attributes of a network traffic follow the ENI which is attached to an EC2 instance or detached from an EC2 instance. When you move an ENI from one EC2 instance to another, network traffic is redirected to the new EC2 instance. You can create and attach additional ENIs to an EC2 instance.
Attaching multiple network interfaces (ENIs) to an EC2 instance is useful to: Create a management network.
Use network and security appliances in your VPC.
Create dual-homed instances with workloads/roles on distinct subnets Create a low-budget, high-availability solution.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.htm|

NEW QUESTION 248
Which one of the below doesn't affect Amazon CloudFront billing?

A. Distribution Type
B. Data Transfer Out
C. Dedicated IP SSL Certificates
D. Requests

Answer: A

Explanation:
Amazon CloudFront is a web service for content delivery. C|oudFront delivers your content using a global network of edge locations and works seamlessly with Amazon S3 which durably stores the original and definitive versions of your files.
Amazon CloudFront billing is maily affected by Data Transfer Out
Edge Location Traffic Distribution Requests
Dedicated IP SSL Certificates
Reference: http://calcu|ator.s3.amazonaws.com/index.htmI

NEW QUESTION 252
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 254
In the most recent company meeting, your CEO focused on the fact that everyone in the organization needs to make sure that all of the infrastructure that is built is truly scalable. Which of the following statements is incorrect in reference to scalable architecture?

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

Answer: C

Explanation:
In AWS it is critical to build a scalable architecture in order to take advantage of a scalable infrastructure. The cloud is designed to provide conceptually infinite scalability. However, you cannot leverage all that scalability in infrastructure if your architecture is not scalable. Both have to work together. You will have to identify the monolithic components and bottlenecks in your architecture, identify the areas where you cannot leverage the on-demand provisioning capabilities in your architecture, and work to refactor your application, in order to leverage the scalable infrastructure and take advantage of the cloud.
Characteristics of a truly scalable application:
Increasing resources results in a proportional increase in performance A scalable service is capable of handling heterogeneity
A scalable service is operationally efficient A scalable service is resilient
A scalable service should become more cost effective when it grows (Cost per unit reduces as the number of units increases)
Reference: http://media.amazonwebservices.com/AWS_Cloud_Best_Practices.pdf

NEW QUESTION 259
A user has defined an AutoScaling termination policy to first delete the instance with the nearest billing hour. AutoScaling has launched 3 instances in the US-East-1A region and 2 instances in the US-East-1 B region. One of the instances in the US-East-1B region is running nearest to the billing hour. Which instance will AutoScaling terminate first while executing the termination action?

A. Random Instance from US-East-1A
B. Instance with the nearest billing hour in US-East-1 B
C. Instance with the nearest billing hour in US-East-1A
D. Random instance from US-East-1B

Answer: C

Explanation:
Even though the user has configured the termination policy, before AutoScaling selects an instance to terminate, it first identifies the Availability Zone that has more instances than the other Availability Zones used by the group. Within the selected Availability Zone, it identifies the instance that matches the specified termination policy.
Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/us-termination-policy.html

NEW QUESTION 264
You are playing around with setting up stacks using JSON templates in CloudFormation to try and understand them a little better. You have set up about 5 or 6 but now start to wonder if you are being charged for these stacks. What is AWS's billing policy regarding stack resources?

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

Answer: B

Explanation:
A stack is a collection of AWS resources that you can manage as a single unit. In other words, you can create, update, or delete a collection of resources by creating, updating, or deleting stacks. All the resources in a stack are defined by the stack's AWS CloudFormation template. A stack, for instance, can include all the resources required to run a web application, such as a web server, a database, and networking rules. If you no longer require that web application, you can simply delete the stack, and all of its related resources are deleted.
You are charged for the stack resources for the time they were operating (even if you deleted the stack right away).
Reference: http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacks.html

NEW QUESTION 265
After deploying a new website for a client on AWS, he asks if you can set it up so that if it fails it can be automatically redirected to a backup website that he has stored on a dedicated server elsewhere. You are wondering whether Amazon Route 53 can do this. Which statement below is correct in regards to Amazon Route 53?

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

Answer: B

Explanation:
With DNS Failover, Amazon Route 53 can help detect an outage of your website and redirect your end users to alternate locations where your application is operating properly.
Reference:
http://aws.amazon.com/about-aws/whats-new/2013/02/11/announcing-dns-failover-for-route-53/

NEW QUESTION 269
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 274
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 276
All Amazon EC2 instances are assigned two IP addresses at launch. Which are those?

A. 2 Elastic IP addresses
B. A private IP address and an Elastic IP address
C. A public IP address and an Elastic IP address
D. A private IP address and a public IP address

Answer: D

Explanation:
In Amazon EC2-Classic every instance is given two IP Addresses: a private IP address and a public IP address
Reference:
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-instance-addressing.html#differences

NEW QUESTION 280
Can I change the EC2 security groups after an instance is launched in EC2-Classic?

A. Yes, you can change security groups after you launch an instance in EC2-Classic.
B. No, you cannot change security groups after you launch an instance in EC2-Classic.
C. Yes, you can only when you remove rules from a security group.
D. Yes, you can only when you add rules to a security grou

Answer: B

Explanation:
After you launch an instance in EC2-Classic, you can't change its security groups. However, you can add rules to or remove rules from a security group, and those changes are automatically applied to all instances that are associated with the security group.
Reference: http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/using-network-security.html

NEW QUESTION 281
George has launched three EC2 instances inside the US-East-1a zone with his AWS account. Ray has launched two EC2 instances in the US-East-Ia zone with his AWS account. Which of the below mentioned statements will help George and Ray understand the availability zone (AZ) concept better?

A. All the instances of George and Ray can communicate over a private IP with a minimal cost
B. The US-East-1a region of George and Ray can be different availability zones
C. All the instances of George and Ray can communicate over a private IP without any cost
D. The instances of George and Ray will be running in the same data centre

Answer: B

Explanation:
Each AWS region has multiple, isolated locations known as Availability Zones. To ensure that the AWS resources are distributed across the Availability Zones for a region, AWS independently maps the Availability Zones to identifiers for each account. In this case the Availability Zone US-East-Ia where George's EC2 instances are running might not be the same location as the US-East-Ia zone of Ray's EC2 instances. There is no way for the user to coordinate the Availability Zones between accounts.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html

NEW QUESTION 286
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 288
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 292
While controlling access to Amazon EC2 resources, which of the following acts as a firewall that controls the traffic allowed to reach one or more instances?

A. A security group
B. An instance type
C. A storage cluster
D. An object

Answer: A

Explanation:
A security group acts as a firewall that controls the traffic allowed to reach one or more instances. When you launch an instance, you assign it one or more security groups.
Reference: http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/UsingIAM.html

NEW QUESTION 297
A user is running a webserver on EC2. The user wants to receive the SMS when the EC2 instance utilization is above the threshold limit. Which AWS services should the user configure in this case?

A. AWS C|oudWatch + AWS SQS.
B. AWS CIoudWatch + AWS SNS.
C. AWS CIoudWatch + AWS SES.
D. AWS EC2 + AWS Cloudwatc

Answer: B

Explanation:

Amazon SNS makes it simple and cost-effective to push to mobile devices, such as iPhone, iPad, Android, Kindle Fire, and internet connected smart devices, as well as pushing to other distributed services. In this case, the user can configure that Cloudwatch sends an alarm on when the threshold is crossed to SNS which will trigger an SMS.
Reference: http://aws.amazon.com/sns/

NEW QUESTION 298
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 299
Your manager has come to you saying that he is very confused about the bills he is receMng from AWS as he is getting different bills for every user and needs you to look into making it more understandable. Which of the following would be the best solution to meet his request?

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

Answer: B

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

NEW QUESTION 303
A user has hosted an application on EC2 instances. The EC2 instances are configured with ELB and Auto Scaling. The application server session time out is 2 hours. The user wants to configure connection draining to ensure that all in-flight requests are supported by ELB even though the instance is being deregistered. What time out period should the user specify for connection draining?

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

Answer: A

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

NEW QUESTION 305
Identify a true statement about the On-Demand instances purchasing option provided by Amazon EC2.

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

Answer: A

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

NEW QUESTION 309
In Amazon EC2, how many Elastic IP addresses can you have by default?

A. 10
B. 2

C. 5
D. 20

Answer: C

Explanation:
The number of Elastic IP addresses you can have in EC2 is 5.
Reference: http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#|imits_ec2

NEW QUESTION 311
You need to create a load balancer in a VPC network that you are building. You can make your load balancer internal (private) or internet-facing (public). When you make your load balancer internal, a DNS name will be created, and it will contain the private IP address of the load balancer. An internal load balancer is not exposed to the internet. When you make your load balancer internet-facing, a DNS name will be created with the public IP address. If you want the Internet-facing load balancer to be connected to the Internet, where must this load balancer reside?

A. The load balancer must reside in a subnet that is connected to the internet using the internet gateway.
B. The load balancer must reside in a subnet that is not connected to the internet.
C. The load balancer must not reside in a subnet that is connected to the internet.
D. The load balancer must be completely outside of your VP

Answer: A

Explanation:
When you create an internal Elastic Load Balancer in a VPC, you need to select private subnets that are in the same Availability Zone as your instances. If the VPC Elastic Load Balancer is to be public facing, you need to create the Elastic Load Balancer in a public subnet. A subnet is a public subnet if it is attached to an Internet Gateway (IGW) with a defined route to that gateway. Selecting more than one public subnet increases the availability of your Elastic Load Balancer.
NB - Elastic Load Balancers in EC2-Classic are always Internet-facing load balancers. Reference:
http://docs.aws.amazon.com/ElasticLoadBalancing/|atest/DeveloperGuide/elb-internet-facing-load-baIan cers.htmI

NEW QUESTION 312
When controlling access to Amazon EC2 resources, each Amazon EBS Snapshot has a attribute that controls which AWS accounts can use the snapshot.

A. createVoIumePermission
B. LaunchPermission
C. SharePermission
D. RequestPermission

Answer: A

Explanation:
Each Amazon EBS Snapshot has a createVoIumePermission attribute that you can set to one or more AWS Account IDs to share the AM with those AWS Accounts. To allow several AWS Accounts to use a particular EBS snapshot, you can use the snapshots's createVoIumePermission attribute to include a list of the accounts that can use it.
Reference: http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/UsingIAM.html

NEW QUESTION 316
Your customer wishes to deploy an enterprise application to AWS which will consist of several web servers, several application servers and a small (50GB) Oracle database information is stored, both in the database and the file systems of the various servers. The backup system must support database recovery whole server and whole disk restores, and indMdual file restores with a recovery time of no more than two hours. They have chosen to use RDS Oracle as the database Which backup architecture will meet these requirements?

A. Backup RDS using automated daily DB backups Backup the EC2 instances using AMs andsupplement with file-level backup to 53 using traditional enterprise backup software to provide fi Ie level restore
B. Backup RDS using a Multi-AZ Deployment Backup the EC2 instances using Amis, and supplement by copying file system data to 53 to provide file level restore.
C. Backup RDS using automated daily DB backups Backup the EC2 instances using EBS snapshots and supplement with file-level backups to Amazon Glacier using traditional enterprise backup software to provide file level restore
D. Backup RDS database to 53 using Oracle RMAN Backup the EC2 instances using Amis, and supplement with EBS snapshots for indMdual volume restore.

Answer: A

Explanation:
Point-In-Time Recovery
In addition to the daily automated backup, Amazon RDS archives database change logs. This enables you to recover your database to any point in time during the backup retention period, up to the last five minutes of database usage.
Amazon RDS stores multiple copies of your data, but for Single-AZ DB instances these copies are stored in a single availability zone. If for any reason a Single-AZ DB instance becomes unusable, you can use point-in-time recovery to launch a new DB instance with the latest restorable data. For more information on working with point-in-time recovery, go to Restoring a DB Instance to a Specified Time.
Note
Mu|ti-AZ deployments store copies of your data in different Availability Zones for greater levels of data durability. For more information on Multi-AZ deployments, see High Availability (MuIti-AZ).

NEW QUESTION 319
......

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