Amazon

Exam Questions AWS-Solution-Architect-Associate

Amazon AWS Certified Solutions Architect - Associate

NEW QUESTION 1
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 2
Can a user get a notification of each instance start / terminate configured with Auto Scaling?

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

Answer: C

Explanation:
The user can get notifications using SNS if he has configured the notifications while creating the Auto Scaling group.
Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/GettingStartedTutorial.html

NEW QUESTION 3
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 4
A client needs you to import some existing infrastructure from a dedicated hosting provider to AWS to try and save on the cost of running his current website. He also needs an automated process that manages backups, software patching, automatic failure detection, and recovery. You are aware that his existing set up currently uses an Oracle database. Which of the following AWS databases would be best for accomplishing this task?

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

Answer: A

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

NEW QUESTION 5
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 6
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 7
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 CIoudHub

Answer: D

Explanation:
If you have multiple VPN connections, you can provide secure communication between sites using the AWS VPN CIoudHub. The VPN CIoudHub 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_CIoudHub.html

NEW QUESTION 8
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 9
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 "EvaIuate 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 10
An Elastic IP address (EIP) is a static IP address designed for dynamic cloud computing. With an EIP, you can mask the failure of an instance or software by rapidly remapping the address to another instance in your account. Your EIP is associated with your AWS account, not a particular EC2 instance, and it remains associated with your account until you choose to explicitly release it. By default how many EIPs is each AWS account limited to on a per region basis?

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

Answer: B

Explanation:
By default, all AWS accounts are limited to 5 Elastic IP addresses per region for each AWS account, because public (IPv4) Internet addresses are a scarce public resource. AWS strongly encourages you to use an EIP primarily for load balancing use cases, and use DNS hostnames for all other inter-node communication.
If you feel your architecture warrants additional EIPs, you would need to complete the Amazon EC2 Elastic IP Address Request Form and give reasons as to your need for additional addresses. Reference:
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-ip-addresses-eip.html#using-instance-ad dressing-limit

NEW QUESTION 10
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 11
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 15
Your EBS volumes do not seem to be performing as expected and your team leader has requested you look into improving their performance. Which of the following is not a true statement relating to the performance of your EBS volumes?

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

Answer: A

Explanation:
Several factors can affect the performance of Amazon EBS volumes, such as instance configuration, I/O characteristics, workload demand, and storage configuration.
Frequent snapshots provide a higher level of data durability, but they may slightly degrade the
performance of your application while the snapshot is in progress. This trade off becomes critical when you have data that changes rapidly. Whenever possible, plan for snapshots to occur during off-peak times in order to minimize workload impact.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSPerformance.htmI

NEW QUESTION 17
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 19
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 22
As AWS grows, most of your clients' main concerns seem to be about security, especially when all of their competitors also seem to be using AWS. One of your clients asks you whether having a competitor who hosts their EC2 instances on the same physical host would make it easier for the competitor to hack into the client's data. Which of the following statements would be the best choice to put your client's mind at rest?

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

Answer: C

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

NEW QUESTION 23
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/AutoScaIing/latest/Deve|operGuide/AS_Concepts.htmI

NEW QUESTION 24
An organization has created an application which is hosted on the AWS EC2 instance. The application stores images to S3 when the end user uploads to it. The organization does not want to store the AWS secure credentials required to access the S3 inside the instance. Which of the below mentioned options is a possible solution to avoid any security threat?

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

Answer: B

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

NEW QUESTION 28
Can resource record sets in a hosted zone have a different domain suffix (for example, www.blog. acme.com and www.acme.ca)?

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

Answer: D

Explanation:
The resource record sets contained in a hosted zone must share the same suffix. For example, the example.com hosted zone can contain resource record sets for www.example.com and wvvw.aws.example.com subdomains, but it cannot contain resource record sets for a www.example.ca subdomain.
Reference: http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/AboutHostedZones.html

NEW QUESTION 31
You have been doing a lot of testing of your VPC Network by deliberately failing EC2 instances to test whether instances are failing over properly. Your customer who will be paying the AWS bill for all this asks you if he being charged for all these instances. You try to explain to him how the billing works on EC2 instances to the best of your knowledge. What would be an appropriate response to give to the customer
in regards to this?

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

Answer: C

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

NEW QUESTION 33
What does the following policy for Amazon EC2 do?
{
"Statement":[{
"Effect":"AI|ow", "Action":"ec2:Describe*", "Resource":"*"
II
}

A. Allow users to use actions that start with "Describe" over all the EC2 resources.
B. Share an AMI with a partner
C. Share an AMI within the account
D. Allow a group to only be able to describe, run, stop, start, and terminate instances

Answer: A

Explanation:
You can use IAM policies to control the actions that your users can perform against your EC2 resources. For instance, a policy with the following statement will allow users to perform actions whose name start with "Describe" against all your EC2 resources.
{
"Statement":[{
"Effect":"AI|ow", "Action":"ec2:Describe*", "Resource":"*"
}I
}
Reference: http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/UsingIAM.htmI

NEW QUESTION 34
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 37
Amazon RDS provides high availability and failover support for DB instances using .

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

Answer: D

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

NEW QUESTION 38
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 41
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 43
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 MuItipart 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 46
You are setting up your first Amazon Virtual Private Cloud (Amazon VPC) so you decide to use the VPC wizard in the AWS console to help make it easier for you.

Which of the following statements is correct regarding instances that you launch into a default subnet via the VPC wizard?

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

Answer: B

Explanation:
Instances that you launch into a default subnet receive both a public IP address and a private IP address. Instances in a default subnet also receive both public and private DNS hostnames. Instances that you launch into a nondefault subnet in a default VPC don't receive a public IP address or a DNS hostname. You can change your subnet's default public IP addressing behavior.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/default-vpc.html

NEW QUESTION 50
Amazon S3 allows you to set per-file permissions to grant read and/or write access. However you have decided that you want an entire bucket with 100 files already in it to be accessible to the public. You don't want to go through 100 files indMdually and set permissions. What would be the best way to do this?

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

Answer: B

Explanation:
Amazon S3 supports several mechanisms that give you filexibility to control who can access your data as well as how, when, and where they can access it. Amazon S3 provides four different access control mechanisms: AWS Identity and Access Management (IAM) policies, Access Control Lists (ACLs), bucket policies, and query string authentication. IAM enables organizations to create and manage multiple users under a single AWS account. With IAM policies, you can grant IAM users fine-grained control to your Amazon S3 bucket or objects. You can use ACLs to selectively add (grant) certain permissions on indMdual objects. Amazon S3 bucket policies can be used to add or deny permissions across some or all of the objects within a single bucket.
With Query string authentication, you have the ability to share Amazon S3 objects through URLs that are
valid for a specified period of time.
Reference: http://aws.amazon.com/s3/detai|s/#security

NEW QUESTION 54
You need to set up a high level of security for an Amazon Relational Database Service (RDS) you have just built in order to protect the confidential information stored in it. What are all the possible security groups that RDS uses?

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

Answer: A

Explanation:
A security group controls the access to a DB instance. It does so by allowing access to IP address ranges or Amazon EC2 instances that you specify.
Amazon RDS uses DB security groups, VPC security groups, and EC2 security groups. In simple terms, a DB security group controls access to a DB instance that is not in a VPC, a VPC security group controls access to a DB instance inside a VPC, and an Amazon EC2 security group controls access to an EC2 instance and can be used with a DB instance.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Welcome.html

NEW QUESTION 57
You have set up an Elastic Load Balancer (ELB) with the usual default settings, which route each request independently to the application instance with the smallest load. However, someone has asked you to bind a user's session to a specific application instance so as to ensure that all requests coming from the user during the session will be sent to the same application instance. AWS has a feature to do this. What is it called?

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

Answer: D

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

NEW QUESTION 58

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 CIoudFront 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 CIoudFront 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 CIoudFront URLs, the CIoudFront 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/AmazonCIoudFront/latest/Deve|operGuide/private-content-restricting-acces s-to-s3.html

NEW QUESTION 61
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 63
Which IAM role do you use to grant AWS Lambda permission to access a DynamoDB Stream?

A. Dynamic role
B. Invocation role
C. Execution role
D. Event Source role

Answer: C

Explanation:
You grant AWS Lambda permission to access a DynamoDB Stream using an IAM role known as the "execution ro|e".
Reference: http://docs.aws.amazon.com/|ambda/latest/dg/intro-permission-model.htm|

NEW QUESTION 64
A scope has been handed to you to set up a super fast gaming server and you decide that you will use Amazon DynamoDB as your database. For efficient access to data in a table, Amazon DynamoDB creates and maintains indexes for the primary key attributes. A secondary index is a data structure that contains a subset of attributes from a table, along with an alternate key to support Query operations. How many types of secondary indexes does DynamoDB support?

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

Answer: A

Explanation:
DynamoDB supports two types of secondary indexes:
Local secondary index — an index that has the same hash key as the table, but a different range key. A local secondary index is "IocaI" in the sense that every partition of a local secondary index is scoped to a table partition that has the same hash key.
Global secondary index — an index with a hash and range key that can be different from those on the table. A global secondary index is considered "gIobaI" because queries on the index can span all of the data in a table, across all partitions.
Reference: http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/SecondaryIndexes.html

NEW QUESTION 65
Select the correct statement: Within Amazon EC2, when using Linux instances, the device name
/dev/sda1 is .

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

Answer: D

Explanation:
Within Amazon EC2, when using a Linux instance, the device name /dev/sda1 is reserved for the root device.

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

NEW QUESTION 66
You need to set up security for your VPC and you know that Amazon VPC provides two features that you can use to increase security for your VPC: security groups and network access control lists (ACLs). You have already looked into security groups and you are now trying to understand ACLs. Which statement below is incorrect in relation to ACLs?

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

Answer: B

Explanation:
Amazon VPC provides two features that you can use to increase security for your VPC:
Security groups—Act as a firewall for associated Amazon EC2 instances, controlling both inbound and outbound traffic at the instance level
Network access control lists (ACLs)—Act as a firewall for associated subnets, controlling both inbound and outbound traffic at the subnet level
Security groups are stateful: (Return traffic is automatically allowed, regardless of any rules) Network ACLs are stateless: (Return traffic must be explicitly allowed by rules)
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Security.html

NEW QUESTION 67
A user comes to you and wants access to Amazon CloudWatch but only wants to monitor a specific LoadBalancer. Is it possible to give him access to a specific set of instances or a specific LoadBalancer?

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

Answer: A

Explanation:
Amazon CloudWatch integrates with AWS Identity and Access Management (IAM) so that you can
specify which CloudWatch actions a user in your AWS Account can perform. For example, you could create an IAM policy that gives only certain users in your organization permission to use GetMetricStatistics. They could then use the action to retrieve data about your cloud resources.
You can't use IAM to control access to CloudWatch data for specific resources. For example, you can't give a user access to CloudWatch data for only a specific set of instances or a specific LoadBalancer. Permissions granted using IAM cover all the cloud resources you use with CloudWatch. In addition, you can't use IAM roles with the Amazon CloudWatch command line tools.
Using Amazon CloudWatch with IAM doesn't change how you use CloudWatch. There are no changes to CloudWatch actions, and no new CloudWatch actions related to users and access control.
Reference: http://docs.aws.amazon.com/AmazonC|oudWatch/latest/DeveloperGuide/UsingIAM.html

NEW QUESTION 72
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 74
You have a lot of data stored in the AWS Storage Gateway and your manager has come to you asking about how the billing is calculated, specifically the Virtual Tape Shelf usage. What would be a correct response to this?

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

Answer: B

Explanation:
The AWS Storage Gateway is a service connecting an on-premises software appliance with cloud-based storage to provide seamless and secure integration between an organization’s on-premises IT environment and AWS’s storage infrastructure.
AWS Storage Gateway billing is as follows. Volume storage usage (per GB per month):
You are billed for the Cached volume data you store in Amazon S3. You are only billed for volume capacity you use, not for the size of the volume you create.

Snapshot Storage usage (per GB per month): You are billed for the snapshots your gateway stores in Amazon S3. These snapshots are stored and billed as Amazon EBS snapshots. Snapshots are incremental backups, reducing your storage charges. When taking a new snapshot, only the data that has changed since your last snapshot is stored.
Virtual Tape Library usage (per GB per month):
You are billed for the virtual tape data you store in Amazon S3. You are only billed for the portion of virtual tape capacity that you use, not for the size of the virtual tape.
Virtual Tape Shelf usage (per GB per month):
You are billed for the virtual tape data you store in Amazon Glacier. You are only billed for the portion of virtual tape capacity that you use, not for the size of the virtual tape.
Reference: https://aws.amazon.com/storagegateway/faqs/

NEW QUESTION 77
You are setting up some EBS volumes for a customer who has requested a setup which includes a RAID (redundant array of inexpensive disks). AWS has some recommendations for RAID setups. Which RAID setup is not recommended for Amazon EBS?

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

Answer: B

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

NEW QUESTION 82
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 84
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 85
The AWS CloudHSM service defines a resource known as a high-availability (HA) ,
which is a virtual partition that represents a group of partitions, typically distributed between several physical HSMs for high-availability.

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

Answer: B

Explanation:
The AWS CloudHSNI service defines a resource known as a high-availability (HA) partition group, which is a virtual partition that represents a group of partitions, typically distributed between several physical HSMs for high-availability.
Reference: http://docs.aws.amazon.com/cloudhsm/latest/userguide/configuring-ha.htmI

NEW QUESTION 90
Is it possible to get a history of all EC2 API calls made on your account for security analysis and operational troubleshooting purposes?

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

Answer: B

Explanation:
To get a history of all EC2 API calls (including VPC and EBS) made on your account, you simply turn on C|oudTrai| in the AWS Management Console.
Reference: https://aws.amazon.com/ec2/faqs/

NEW QUESTION 93
After moving an E-Commerce website for a client from a dedicated server to AWS you have also set up auto scaling to perform health checks on the instances in your group and replace instances that fail these checks. Your client has come to you with his own health check system that he wants you to use as it has proved to be very useful prior to his site running on AWS. What do you think would be an appropriate response to this given all that you know about auto scaling?

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

Answer: C

Explanation:
Auto Scaling periodically performs health checks on the instances in your group and replaces instances that fail these checks. By default, these health checks use the results of EC2 instance status checks to determine the health of an instance. If you use a load balancer with your Auto Scaling group, you can optionally choose to include the results of Elastic Load Balancing health checks.
Auto Scaling marks an instance unhealthy if the calls to the Amazon EC2 action DescribeInstanceStatus returns any other state other than running, the system status shows impaired, or the calls to Elastic Load Balancing action DescribeInstanceHealth returns OutOfService in the instance state field.
After an instance is marked unhealthy because of an Amazon EC2 or Elastic Load Balancing health check, it is scheduled for replacement.
You can customize the health check conducted by your Auto Scaling group by specifying additional checks or by having your own health check system and then sending the instance's health information directly from your system to Auto Scaling.
Reference: http://docs.aws.amazon.com/AutoScaling/latest/Deve|operGuide/healthcheck.html

NEW QUESTION 95
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 98
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 101
You are using Amazon SES as an email solution but are unsure of what its limitations are. Which statement below is correct in regards to that?

A. New Amazon SES users who have received production access can send up to 1,000 emails per 24-hour period, at a maximum rate of 10 emails per second.
B. Every Amazon SES sender has a the same set of sending limits

C. Sending limits are based on messages rather than on recipients
D. Every Amazon SES sender has a unique set of sending limits

Answer: D

Explanation:
Amazon Simple Email Service (Amazon SES) is a highly scalable and cost-effective email-sending
service for businesses and developers. Amazon SES eliminates the complexity and expense of building an in-house email solution or licensing, installing, and operating a third-party email service for this type of email communication.
Every Amazon SES sender has a unique set of sending limits, which are calculated by Amazon SES on an ongoing basis:
Sending quota — the maximum number of emails you can send in a 24-hour period. Maximum send rate — the maximum number of emails you can send per second.
New Amazon SES users who have received production access can send up to 10,000 emails per 24-hour period, at a maximum rate of 5 emails per second.
Amazon SES automatically adjusts these limits upward, as long as you send high-quality email. If your existing quota is not adequate for your needs and the system has not automatically increased your quota, you can submit an SES Sending Quota Increase case at any time.
Sending limits are based on recipients ratherthan on messages. You can check your sending limits at any time by using the Amazon SES console.
Note that if your email is detected to be of poor or QUESTION able quality (e.g., high complaint rates, high bounce rates, spam, or abusive content), Amazon SES might temporarily or permanently reduce your permitted send volume, or take other action as AWS deems appropriate.
Reference: https://aws.amazon.com/ses/faqs/

NEW QUESTION 102
You have set up an S3 bucket with a number of images in it and you have decided that you want anybody to be able to access these images, even anonymous users. To accomplish this you create a bucket policy. You will need to use an Amazon S3 bucket policy that specifies a in the principal element, which means anyone can access the bucket.

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

Answer: C

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

NEW QUESTION 105
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 C|oudFront distribution, specify the URL of the load balancer for the domain name of your origin server.
Reference: http://docs.aws.amazon.com/AmazonC|oudFront/latest/DeveIoperGuide/CustomOriginBestPractices.htmI

NEW QUESTION 110
What is the time period with which metric data is sent to CIoudWatch when detailed monitoring is enabled on an Amazon EC2 instance?

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

Answer: C

Explanation:
By default, Amazon EC2 metric data is automatically sent to CIoudWatch in 5-minute periods. However, you can, enable detailed monitoring on an Amazon EC2 instance, which sends data to CIoudWatch in
1-minute periods
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-cloudwatch.htmI

NEW QUESTION 114
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 116
What happens to data on an ephemeral volume of an EBS-backed EC2 instance if it is terminated or if it fails?

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

Answer: D

Explanation:
Any data on the instance store volumes persists as long as the instance is running, but this data is deleted when the instance is terminated or if it fails (such as if an underlying drive has issues). After an instance store-backed instance fails or terminates, it cannot be restored.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/RootDeviceStorage.html

NEW QUESTION 121
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 122
A company wants to review the security requirements of Glacier. Which of the below mentioned statements is true with respect to the AWS Glacier data security?

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

Answer: A

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

NEW QUESTION 124
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 126
Your manager has asked you to set up a public subnet with instances that can send and receive internet traffic, and a private subnet that can't receive traffic directly from the internet, but can initiate traffic to the internet (and receive responses) through a NAT instance in the public subnet. Hence, the following 3 rules need to be allowed:
Inbound SSH traffic.
Web sewers in the public subnet to read and write to MS SQL servers in the private subnet Inbound RDP traffic from the Microsoft Terminal Services gateway in the public private subnet What are the respective ports that need to be opened for this?

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

Answer: A

Explanation:
A network access control list (ACL) is an optional layer of security that acts as a firewall for controlling traffic in and out of a subnet. You might set up network ACLs with rules similar to your security groups in order to add an additional layer of security to your VPC.
The following ports are recommended by AWS for a single subnet with instances that can receive and send Internet traffic and a private subnet that can't receive traffic directly from the Internet. However, it can initiate traffic to the Internet (and receive responses) through a NAT instance in the public subnet. Inbound SSH traffic. Port 22
Web sewers in the public subnet to read and write to MS SQL sewers in the private subnet. Port 1433 Inbound RDP traffic from the Microsoft Terminal Sewices gateway in the public private subnet. Port 3389 Reference:
http://docs.aws.amazon.com/AmazonVPC/Iatest/UserGuide/VPC_Appendix_NACLs.htm|#VPC_Appendi x_NAC Ls_Scenario_2

NEW QUESTION 131
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 133
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 138
A user has created an ELB with Auto Scaling. Which of the below mentioned offerings from ELB helps the
user to stop sending new requests traffic from the load balancer to the EC2 instance when the instance is being deregistered while continuing in-flight requests?

A. ELB sticky session
B. ELB deregistration check
C. ELB auto registration Off
D. ELB connection draining

Answer: D

Explanation:
The Elastic Load Balancer connection draining feature causes the load balancer to stop sending new requests to the back-end instances when the instances are deregistering or become unhealthy, while ensuring that in-flight requests continue to be served.
Reference:
http://docs.aws.amazon.com/ElasticLoadBaIancing/latest/DeveIoperGuide/config-conn-drain.htmI

NEW QUESTION 141
You have some very sensitive data stored on AWS S3 and want to try every possible alternative to keeping it secure in regards to access control. What are the mechanisms available for access control on AWS S3?

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

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

Answer: A

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

NEW QUESTION 142
You have been asked to set up a database in AWS that will require frequent and granular updates. You know that you will require a reasonable amount of storage space but are not sure of the best option. What is the recommended storage option when you run a database on an instance with the above criteria?

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

Answer: B

Explanation:
Amazon EBS provides durable, block-level storage volumes that you can attach to a running Amazon EC2 instance. You can use Amazon EBS as a primary storage device for data that requires frequent and granular updates. For example, Amazon EBS is the recommended storage option when you run a database on an instance.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Storage.html

NEW QUESTION 144
You have been asked to set up monitoring of your network and you have decided that Cloudwatch would be the best service to use. Amazon CIoudWatch monitors your Amazon Web Services (AWS) resources and the applications you run on AWS in real-time. You can use CIoudWatch to collect and track metrics, which are the variables you want to measure for your resources and applications. Which of the following items listed can AWS Cloudwatch monitor?

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

Answer: B

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

NEW QUESTION 148
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/generaI/latest/gr/aws_service_limits.htmI#|imits_ec2

NEW QUESTION 150
Can you move a Reserved Instance from one Availability Zone to another?

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

Answer: A

Explanation:
Each Reserved Instance is associated with a specific Region, which is fixed for the lifetime of the reservation and cannot be changed. Each reservation can, however, be used in any of the available AZs within the associated Region.
Reference: https://aws.amazon.com/rds/faqs/

NEW QUESTION 153
When controlling access to Amazon EC2 resources, each Amazon EBS Snapshot has a attribute that controls which AWS accounts can use the snapshot.

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

Answer: A

Explanation:
Each Amazon EBS Snapshot has a createVolumePermission 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 createVolumePermission attribute to include a list of the accounts that can use it.
Reference: http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/UsingIAM.html

NEW QUESTION 154
Your company has HQ in Tokyo and branch offices all over the world and is using a logistics software with a multi-regional deployment on AWS in Japan, Europe and USA, The logistic software has a 3- tier architecture and currently uses MySQL 5.6 for data persistence. Each region has deployed its own database
In the HQ region you run an hourly batch process reading data from every region to compute cross regional reports that are sent by email to all offices this batch process must be completed as fast as possible to quickly optimize logistics how do you build the database architecture in order to meet the requirements'?

A. For each regional deployment, use RDS MySQL with a master in the region and a read replica in the HQ region
B. For each regional deployment, use MySQL on EC2 with a master in the region and send hourly EBS snapshots to the HQ region
C. For each regional deployment, use RDS MySQL with a master in the region and send hourly RDS snapshots to the HQ region
D. For each regional deployment, use MySQL on EC2 with a master in the region and use 53 to copy data files hourly to the HQ region
E. Use Direct Connect to connect all regional MySQL deployments to the HQ region and reduce network latency for the batch process

Answer: A

NEW QUESTION 158
You have recently joined a startup company building sensors to measure street noise and air quality in urban areas. The company has been running a pilot deployment of around 100 sensors for 3 months each sensor uploads 1KB of sensor data every minute to a backend hosted on AWS.
During the pilot, you measured a peak or 10 IOPS on the database, and you stored an average of 3GB of sensor data per month in the database.
The current deployment consists of a load-balanced auto scaled Ingestion layer using EC2 instances and a PostgreSQL RDS database with 500GB standard storage.
The pilot is considered a success and your CEO has managed to get the attention or some potential investors. The business plan requires a deployment of at least IOOK sensors which needs to be supported by the backend. You also need to store sensor data for at least two years to be able to compare year over year Improvements.
To secure funding, you have to make sure that the platform meets these requirements and leaves room for further scaling. Which setup win meet the requirements?

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

Answer: C

NEW QUESTION 163
Your company runs a customer facing event registration site This site is built with a 3-tier architecture with web and application tier servers and a MySQL database The application requires 6 web tier servers and 6 application tier servers for normal operation, but can run on a minimum of 65% server capacity and a single MySQL database. When deploying this application in a region with three availability zones (AZs) which architecture provides high availability?

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

Answer: D

Explanation:
Amazon RDS MuIti-AZ Deployments
Amazon RDS Multi-AZ deployments provide enhanced availability and durability for Database (DB) Instances, making them a natural fit for production database workloads. When you provision a MuIti-AZ DB Instance, Amazon RDS automatically creates a primary DB Instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). Each AZ runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. In case of an infrastructure failure (for example, instance hardware failure, storage failure, or network disruption), Amazon RDS performs an automatic failover to the standby, so that you can resume database operations as soon as the failover is complete. Since the endpoint for your DB Instance remains the same after a failover, your application can resume database operation without the need for manual administrative intervention.
Enhanced Durability
MuIti-AZ deployments for the MySQL, Oracle, and PostgreSQL engines utilize synchronous physical replication to keep data on the standby up-to-date with the primary. MuIti-AZ deployments for the SQL Server engine use synchronous logical replication to achieve the same result, employing SQL

Server-native Mrroring technology. Both approaches safeguard your data in the event of a DB Instance failure or loss of an Availability Zone.
If a storage volume on your primary fails in a Multi-AZ deployment, Amazon RDS automatically initiates a failover to the up-to-date standby. Compare this to a Single-AZ deployment: in case of a Single-AZ database failure, a user-initiated point-in-time-restore operation will be required. This operation can take several hours to complete, and any data updates that occurred after the latest restorable time (typically within the last five minutes) will not be available.
Amazon Aurora employs a highly durable, SSD-backed virtualized storage layer purpose-built for database workloads. Amazon Aurora automatically replicates your volume six ways, across three Availability Zones. Amazon Aurora storage is fault-tolerant, transparently handling the loss of up to two copies of data without affecting database write availability and up to three copies without affecting read availability. Amazon Aurora storage is also self-healing. Data blocks and disks are continuously scanned for errors and replaced automatically.
Increased Availability
You also benefit from enhanced database availability when running Multi-AZ deployments. If an Availability Zone failure or DB Instance failure occurs, your availability impact is limited to the time automatic failover takes to complete: typically under one minute for Amazon Aurora and one to two minutes for other database engines (see the RDS FAQ for details).
The availability benefits of MuIti-AZ deployments also extend to planned maintenance and backups. In the case of system upgrades like QS patching or DB Instance scaling, these operations are applied first on
the standby, prior to the automatic failover. As a result, your availability impact is, again, only the time required for automatic failover to complete.
Unlike Single-AZ deployments, 1/0 actMty is not suspended on your primary during backup for MuIti-AZ deployments for the MySQL, Oracle, and PostgreSQL engines, because the backup is taken from the standby. However, note that you may still experience elevated latencies for a few minutes during backups for MuIti-AZ deployments.
On instance failure in Amazon Aurora deployments, Amazon RDS uses RDS MuIti-AZ technology to automate failover to one of up to 15 Amazon Aurora Replicas you have created in any of three Availability Zones. If no Amazon Aurora Replicas have been provisioned, in the case of a failure, Amazon RDS will attempt to create a new Amazon Aurora DB instance for you automatically.

NEW QUESTION 168
Your application is using an ELB in front of an Auto Scaling group of web/application sewers deployed across two AZs and a MuIti-AZ RDS Instance for data persistence.
The database CPU is often above 80% usage and 90% of 1/0 operations on the database are reads. To improve performance you recently added a single-node Memcached EIastiCache Cluster to cache frequent DB query results. In the next weeks the overall workload is expected to grow by 30%.
Do you need to change anything in the architecture to maintain the high availability or the application with the anticipated additional load? Why?

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

Answer: A

Explanation:
EIastiCache for Memcached
The primary goal of caching is typically to offload reads from your database or other primary data source. In most apps, you have hot spots of data that are regularly queried, but only updated periodically. Think of the front page of a blog or news site, or the top 100 leaderboard in an online game. In this type of case, your app can receive dozens, hundreds, or even thousands of requests for the same data before it's updated again. Having your caching layer handle these queries has several advantages. First, it's considerably cheaper to add an in-memory cache than to scale up to a larger database cluster. Second,
an in-memory cache is also easier to scale out, because it's easier to distribute an in-memory cache horizontally than a relational database.
Last, a caching layer provides a request buffer in the event of a sudden spike in usage. If your app or game ends up on the front page of Reddit or the App Store, it's not unheard of to see a spike that is 10 to 100 times your normal application load. Even if you autoscale your application instances, a IOx request spike will likely make your database very unhappy.
Let's focus on EIastiCache for Memcached first, because it is the best fit for a caching focused solution. We'II revisit Redislater in the paper, and weigh its advantages and disadvantages.
Architecture with EIastiCache for Memcached
When you deploy an EIastiCache Memcached cluster, it sits in your application as a separate tier alongside your database. As mentioned previously, Amazon EIastiCache does not directly communicate with your database tier, or indeed have any particular knowledge of your database. A simplified deployment for a web application looks something like this:

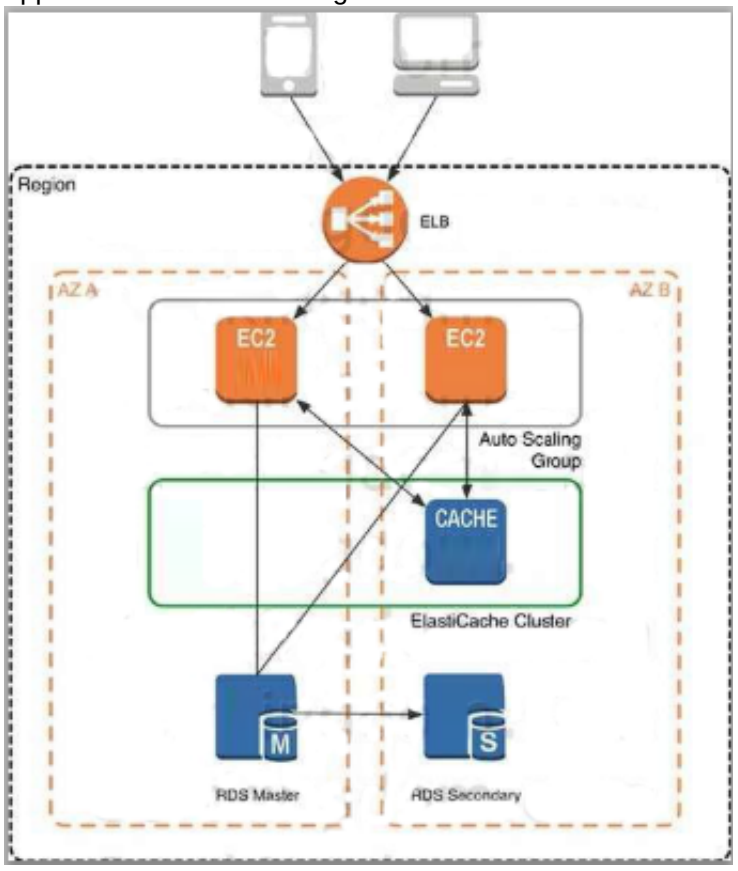

In this architecture diagram, the Amazon EC2 application instances are in an Auto Scaling group, located behind a load balancer using Elastic Load Balancing, which distributes requests among the instances. As requests come into a given EC2 instance, that EC2 instance is responsible for communicating with EIastiCache and the database tier. For development purposes, you can begin with a single EIastiCache node to test your application, and then scale to additional

cluster nodes by modifying t he EIastiCache cluster. As you add additional cache nodes, the EC2 application instances are able to distribute cache keys across multiple EIastiCache nodes. The most common practice is to use client-side sharding to distribute keys across cache nodes, which we will discuss later in this paper.

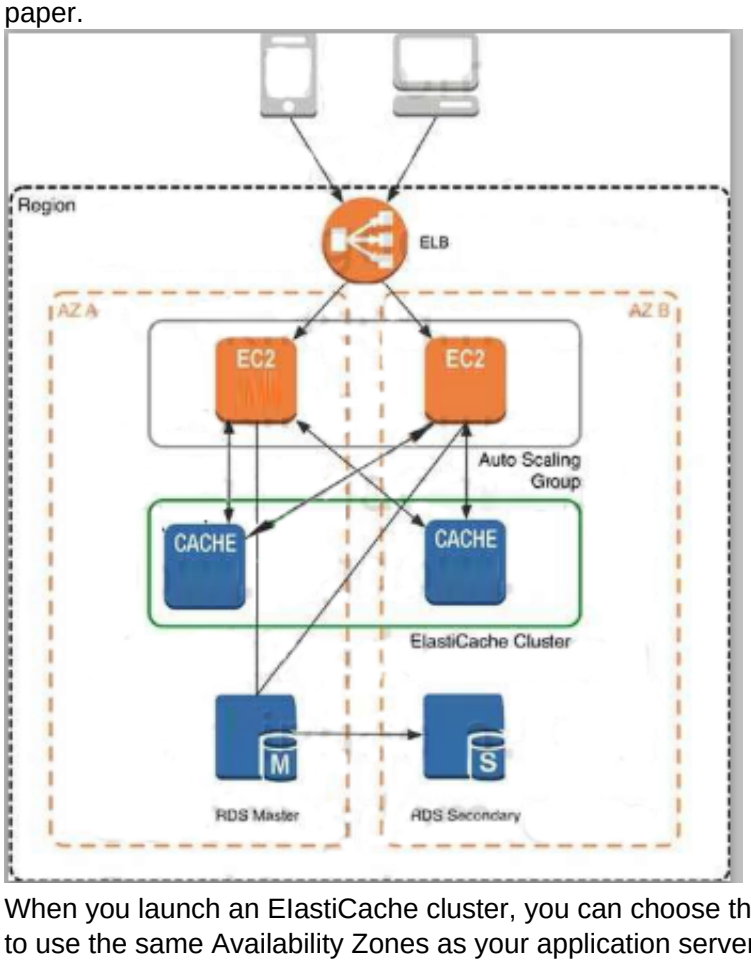

When you launch an EIastiCache cluster, you can choose the Availability Zone(s) that the cluster lives in. For best performance, you should configure your cluster to use the same Availability Zones as your application servers. To launch an EIastiCache cluster in a specific Availability Zone, make sure to specify the Preferred Zone(s) option during cache cluster creation. The Availability Zones that you specify will be where EIastiCache will launch your cache nodes. We recommend that you select Spread Nodes Across Zones, which tells EIastiCache to distribute cache nodes across these zones as evenly as possible. This distribution will mitigate the impact of an Availability Zone disruption on your E|astiCache nodes. The trade-off is that some of the requests from your application to EIastiCache will go to a node in a different Availability Zone, meaning latency will be slightly higher.
For more details, refer to Creating a Cache Cluster in the Amazon EIastiCache User Guide.
As mentioned at the outset, EIastiCache can be coupled with a wide variety of databases. Here is an example architecture that uses Amazon DynamoDB instead of Amazon RDS and IV|ySQL:

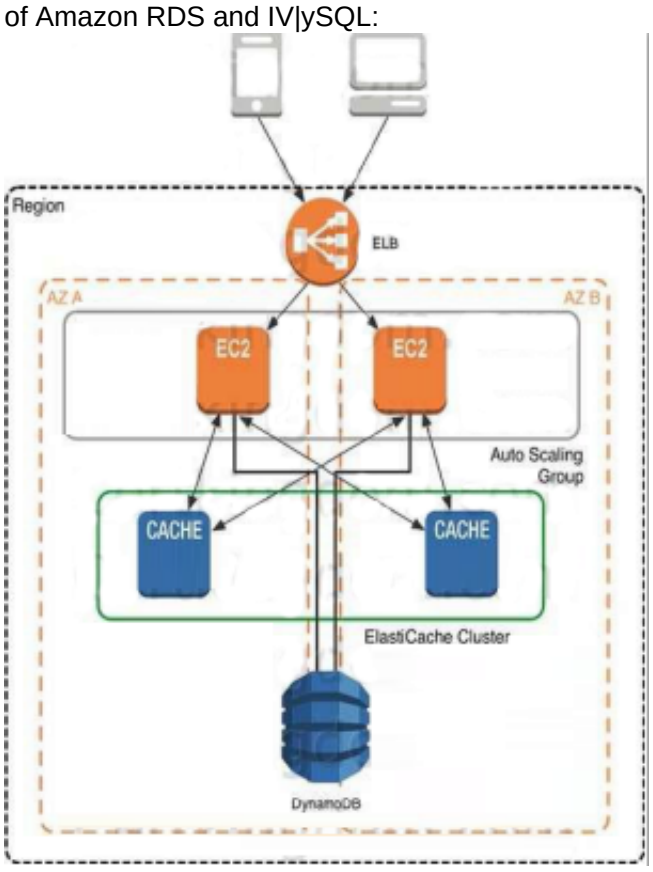

This combination of DynamoDB and EIastiCache is very popular with mobile and game companies, because DynamoDB allows for higher write throughput at lower cost than traditional relational databases. In addition, DynamoDB uses a key-value access pattern similar to EIastiCache, which also simplifies the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

NEW QUESTION 173
Refer to the architecture diagram above of a batch processing solution using Simple Queue Service (SQS) to set up a message queue between EC2 instances which are used as batch processors Cloud Watch monitors the number of Job requests (queued messages) and an Auto Scaling group adds or deletes batch sewers automatically based on parameters set in Cloud Watch alarms. You can use this architecture to implement which of the following features in a cost effective and efficient manner?

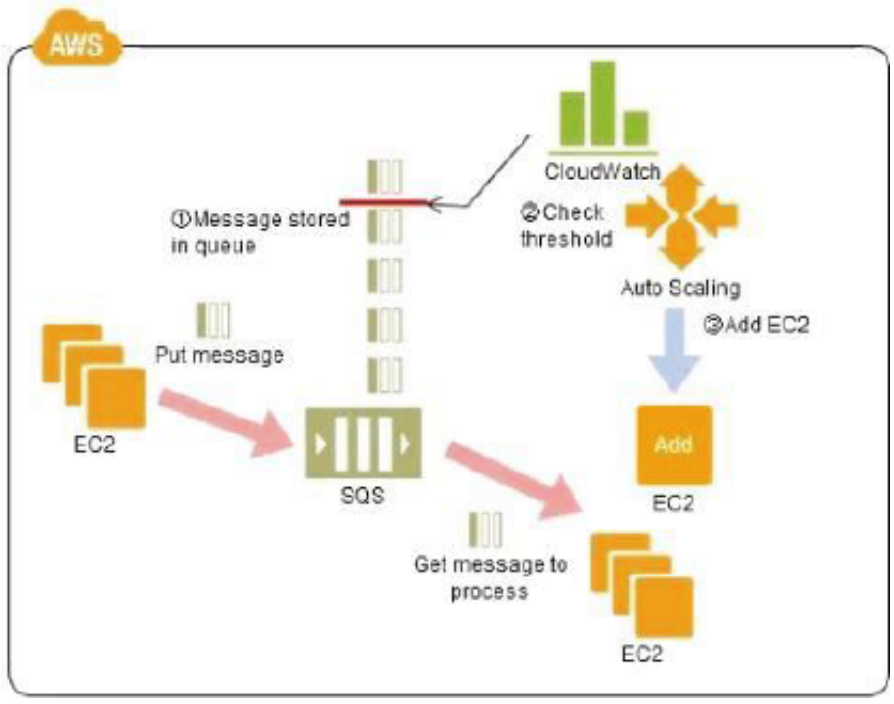

A. Reduce the overall lime for executing jobs through parallel processing by allowing a busy EC2 instance that receives a message to pass it to the next instance in a daisy-chain setup.
B. Implement fault tolerance against EC2 instance failure since messages would remain in SQS and worn can continue with recovery of EC2 instances implement fault tolerance against SQS failure by backing up messages to 53.
C. Implement message passing between EC2 instances within a batch by exchanging messages through SQS.
D. Coordinate number of EC2 instances with number of job requests automatically thus Improving cost effectiveness.
E. Handle high priority jobs before lower priority jobs by assigning a priority metadata fie Id to SQS messages.

Answer: D

Explanation:
Reference:
There are cases where a large number of batch jobs may need processing, and where the jobs may need to be re-prioritized.
For example, one such case is one where there are differences between different levels of services for unpaid users versus subscriber users (such as the time until publication) in services enabling, for example, presentation fi les to be uploaded for publication from a web browser. When the user uploads a presentation file, the conversion processes, for example, for publication are performed as batch
processes on the system side, and the file is published after the conversion. Is it then necessary to be able to assign the level of priority to the batch processes for each type of subscriber.
Explanation of the Cloud Solution/Pattern
A queue is used in controlling batch jobs. The queue need only be provided with priority numbers. Job requests are controlled by the queue, and the job requests in the queue are processed by a batch server. In Cloud computing, a highly reliable queue is provided as a service, which you can use to
structure a highly reliable batch system with ease. You may prepare multiple queues depending on priority levels, with job requests put into the queues depending on their priority levels, to apply prioritization to batch processes. The performance (number) of batch servers corresponding to a queue must be in accordance with the priority level thereof.
Implementation
In AWS, the queue service is the Simple Queue Service (SQS). MuItipIe SQS queues may be prepared to prepare queues for indMdual priority levels (with a priority queue and a secondary queue).
Moreover, you may also use the message Delayed Send function to delay process execution. Use SQS to prepare multiple queues for the indMdual priority levels. Place those processes to be executed immediately (job requests) in the high priority queue. Prepare numbers of batch servers, for processing the job requests of the queues, depending on the priority levels.
Queues have a message "Delayed Send" function. You can use this to delay the time for starting a process.
Configuration
Benefits
You can increase or decrease the number of servers for processing jobs to change automatically the processing speeds of the priority queues and secondary queues.
You can handle performance and service requirements through merely increasing or decreasing the number of EC2 instances used in job processing.
Even if an EC2 were to fail, the messages (jobs) would remain in the queue service, enabling processing to be continued immediately upon recovery of the EC2 instance, producing a system that is robust to failure.
Cautions
Depending on the balance between the number of EC2 instances for performing the processes and the number of messages that are queued, there may be cases where processing in the secondary queue may be completed first, so you need to monitor the processing speeds in the primary queue and the secondary queue.

NEW QUESTION 178
Your company currently has a 2-tier web application running in an on-premises data center. You have experienced several infrastructure failures in the past two months resu Iting in significant financial losses. Your CIO is strongly agreeing to move the application to AWS. While working on achieving buy-in from the other company executives, he asks you to develop a disaster recovery plan to help improve Business continuity in the short term. He specifies a target Recovery Time Objective (RTO) of 4 hours and a Recovery Point Objective (RPO) of 1 hour or less. He also asks you to implement the solution within 2 weeks. Your database is 200GB in size and you have a 20Mbps Internet connection.
How would you do this while minimizing costs?

A. Create an EBS backed private AMI which includes a fresh install of your applicatio
B. Develop a CIoudFormation template which includes your AMI and the required EC2, AutoScaIing, and ELB resources to support deploying the application across Multiple- Availability-Zone
C. Asynchronously replicate transactions from your on-premises database to a database instance in AWS across a secure VPN connection.
D. Deploy your application on EC2 instances within an Auto Scaling group across multiple availability zone
E. Asynchronously replicate transactions from your on-premises database to a database instance in AWS across a secure VPN connection.
F. Create an EBS backed private AMI which includes a fresh install of your applicatio
G. Setup a script in your data center to backup the local database every 1 hour and to encrypt and copy the resulting file to an 53 bucket using multi-part upload.
H. Install your application on a compute-optimized EC2 instance capable of supporting the application 's average loa
I. Synchronously replicate transactions from your on-premises database to a database instance in AWS across a secure Direct Connect connection.

Answer: A

Explanation:
Overview of Creating Amazon EBS-Backed AMIs
First, launch an instance from an AMI that's similar to the AMI that you'd like to create. You can connect to your instance and customize it. When the instance is configured correctly, ensure data integrity by
stopping the instance before you create an AMI, then create the image. When you create an Amazon EBS-backed AMI, we automatically register it for you.
Amazon EC2 powers down the instance before creating the AMI to ensure that everything on the instance is stopped and in a consistent state during the creation process. If you're confident that your instance is in a consistent state appropriate for AMI creation, you can tell Amazon EC2 not to power down and reboot the instance. Some file systems, such as XFS, can freeze and unfreeze actMty, making it safe to create the image without rebooting the instance.
During the AMI-creation process, Amazon EC2 creates snapshots of your instance's root volume and any other EBS volumes attached to your instance. If any volumes attached to the instance are encrypted, the new AMI only launches successfully on instances that support Amazon EBS encryption. For more information, see Amazon EBS Encryption.
Depending on the size of the volumes, it can take several minutes for the AMI-creation process to complete (sometimes up to 24 hours).You may find it more efficient to create snapshots of your volumes prior to creating your AMI. This way, only small, incremental snapshots need to be created when the AMI is created, and the process completes more quickly (the total time for snapshot creation remains the same). For more information, see Creating an Amazon EBS Snapshot.
After the process completes, you have a new AMI and snapshot created from the root volume of the instance. When you launch an instance using the new AMI, we create a new EBS volume for its root volume using the snapshot. Both the AMI and the snapshot incur charges to your account until you delete them. For more information, see Deregistering Your AMI.
If you add instance-store volumes or EBS volumes to your instance in addition to the root device volume, the block device mapping for the new AMI contains information for these volumes, and the block device mappings for instances that you launch from the new AMI automatically contain information for these volumes. The instance-store volumes specified in the block device mapping for the new instance are new and don't contain any data from the instance store volumes of the instance you used to create the AMI. The data on EBS volumes persists. For more information, see Block Device Mapping.

NEW QUESTION 183
Your company hosts a social media site supporting users in multiple countries. You have been asked to provide a highly available design tor the application that leverages multiple regions tor the most recently accessed content and latency sensitive portions of the wet) site The most latency sensitive component of the application involves reading user preferences to support web site personalization and ad selection. In addition to running your application in multiple regions, which option will support this app|ication's requirements?

A. Serve user content from 53. CIoudFront and use Route53 latency-based routing between ELBs in each region Retrieve user preferences from a local DynamoDB table in each region and leverage SQS to capture changes to user preferences with 505 workers for propagating updates to each table.
B. Use the 53 Copy API to copy recently accessed content to multiple regions and serve user content from 53. C|oudFront with dynamic content and an ELB in each region Retrieve user preferences from an EIasticCache cluster in each region and leverage SNS notifications to propagate user preference changes to a worker node in each region.
C. Use the 53 Copy API to copy recently accessed content to multiple regions and serve user content from 53 CIoudFront and Route53 latency-based routing Between ELBs In each region Retrieve user preferences from a DynamoDB table and leverage SQS to capture changes to user preferences with 505 workers for propagating DynamoDB updates.
D. Serve user content from 53. CIoudFront with dynamic content, and an ELB in each region Retrieve user preferences from an EIastiCache cluster in each region and leverage Simple Workflow (SWF) to manage the propagation of user preferences from a centralized OB to each EIastiCache cluster.

Answer: A

NEW QUESTION 187
A customer has established an AWS Direct Connect connection to AWS. The link is up and routes are being advertised from the customer's end, however the customer is unable to connect from EC2 instances inside its VPC to servers residing in its datacenter.
Which of the following options provide a viable solution to remedy this situation? (Choose 2 answers)

A. Add a route to the route table with an IPsec VPN connection as the target.
B. Enable route propagation to the virtual pinnate gateway (VGW).
C. Enable route propagation to the customer gateway (CGW).
D. Modify the route table of all Instances using the 'route' command.
E. Modify the Instances VPC subnet route table by adding a route back to the customer's on-premises environment.

Answer: AC

NEW QUESTION 189
You are implementing AWS Direct Connect. You intend to use AWS public service end points such as Amazon 53, across the AWS Direct Connect link. You want other Internet traffic to use your existing link to an Internet Service Provider.
What is the correct way to configure AW5 Direct connect for access to services such as Amazon 53?

A. Configure a public Interface on your AW5 Direct Connect link Configure a static route via your AW5 Direct Connect link that points to Amazon 53 Advertise a default route to AW5 using BGP.
B. Create a private interface on your AW5 Direct Connect lin
C. Configure a static route via your AW5 Direct connect link that points to Amazon 53 Configure specific routes to your network in your VPC,
D. Create a public interface on your AW5 Direct Connect link Redistribute BGP routes into your existing routing infrastructure advertise specific routes for your network to AW5.
E. Create a private interface on your AW5 Direct connect lin
F. Redistribute BGP routes into your existing routing infrastructure and advertise a default route to AW5.

Answer: C

NEW QUESTION 194
You are migrating a legacy client-server application to AWS. The application responds to a specific DNS domain (e.g. www.examp|e.com) and has a 2-tier architecture, with multiple application sewers and a database server. Remote clients use TCP to connect to the application sewers. The application servers need to know the IP address of the clients in order to function properly and are currently taking that information from the TCP socket. A MuIti-AZ RDS MySQL instance will be used for the database. During the migration you can change the application code, but you have to file a change request.
How would you implement the architecture on AWS in order to maximize scalability and high availability?

A. File a change request to implement Alias Resource support in the applicatio
B. Use Route 53 Alias Resource Record to distribute load on two application servers in different AZs.
C. File a change request to implement Latency Based Routing support in the applicatio
D. Use Route 53 with Latency Based Routing enabled to distribute load on two application servers in different AZs.
E. File a change request to implement Cross-Zone support in the applicatio
F. Use an ELB with a TCP Listener and Cross-Zone Load Balancing enabled, two application servers in different AZs.
G. File a change request to implement Proxy Protocol support in the applicatio
H. Use an ELB with a TCP Listener and Proxy Protocol enabled to distribute load on two application servers in different AZs.

Answer: D

NEW QUESTION 196
You require the ability to analyze a large amount of data, which is stored on Amazon 53 using Amazon Elastic Map Reduce. You are using the cc2 8x large Instance type, whose CPUs are mostly idle during processing. Which of the below would be the most cost efficient way to reduce the runtime of the job?

A. Create more smaller flies on Amazon 53.
B. Add additional cc2 8x large instances by introducing a task group.
C. Use smaller instances that have higher aggregate 1/0 performance.
D. Create fewer, larger fi les on Amazon 53.

Answer: C

NEW QUESTION 197
Your department creates regular analytics reports from your company's log files All log data is collected in Amazon 53 and processed by daily Amazon Elastic MapReduce (EMR) jobs that generate daily PDF reports and aggregated tables in CSV format for an Amazon Redshift data warehouse.
Your CFO requests that you optimize the cost structure for this system.
Which of the following alternatives will lower costs without compromising average performance of the system or data integrity for the raw data?

A. Use reduced redundancy storage (RRS) for all data In 53. Use a combination of Spot Instances and Reserved Instances for Amazon EMR job
B. Use Reserved Instances for Amazon Redshift.
C. Use reduced redundancy storage (RRS) for PDF and .csv data in 53. Add Spot Instances to EMR job
D. Use Spot Instances for Amazon Redshift.
E. Use reduced redundancy storage (RRS) for PDF and .csv data In Amazon 53. Add Spot Instances to Amazon EMR job
F. Use Reserved Instances for Amazon Redshift.
G. Use reduced redundancy storage (RRS) for all data in Amazon 53. Add Spot Instances to Amazon EMR job
H. Use Reserved Instances for Amazon Redshift.

Answer: C

Explanation:
Using Reduced Redundancy Storage
Amazon 53 stores objects according to their storage class. It assigns the storage class to an object when it is written to Amazon 53. You can assign objects a specific sto rage class (standard or reduced redundancy) only when you write the objects to an Amazon 53 bucket or when you copy objects that are already stored in Amazon 53. Standard is the default storage class. For information about storage classes, see Object Key and Metadata.
In order to reduce storage costs, you can use reduced redundancy storage for noncritical, reproducible data at lower levels of redundancy than Amazon 53 provides with standard storage. The lower level of redundancy results in less durability and availability, but in many cases, the lower costs can make reduced redundancy storage an acceptable storage solution. For example, it can be a cost effective solution for sharing media content that is durably stored elsewhere. It can also make sense if you are storing thumbnails and other resized images that can be easily reproduced from an original image. Reduced redundancy storage is designed to provide 99.99% durability of objects over a given year.
This durability level corresponds to an average annual expected loss of 0.01% of objects. For example, if you store 10,000 objects using the RRS option, you can, on average, expect to incur an annual loss of a single object per year (0.01% of 10,000 objects).
Note
This annual loss represents an expected average and does not guarantee the loss of less than 0.01% of objects in a given year.
Reduced redundancy storage stores objects on multiple devices across multiple facilities, providing 400 times the durability of a typical disk drive, but it does not replicate objects as many times as Amazon 53 standard storage. In addition, reduced redundancy storage is designed to sustain the loss of data in a single facility.
If an object in reduced redundancy storage has been lost, Amazon 53 will return a 405 error on requests made to that object. Amazon 53 also offers notifications for reduced redundancy storage object loss: you can configure your bucket so that when Amazon 53 detects the loss of an RRS object, a notification will be sent through Amazon Simple Notification Service (Amazon SNS). You can then replace the lost object. To enable notifications, you can use the Amazon 53 console to set the Notifications property of your bucket.

NEW QUESTION 199
An AWS customer is deploying an application mat is composed of an AutoScaling group of EC2 Instances.
The customers security policy requires that every outbound connection from these instances to any other service within the customers Virtual Private Cloud must be authenticated using a unique x 509 certificate that contains the specific instance-id.
In addition an x 509 certificates must Designed by the customer's Key management service in order to be trusted for authentication.
Which of the following configurations will support these requirements?

A. Configure an IAM Role that grants access to an Amazon 53 object containing a signed certificate and configure me Auto Scaling group to launch instances with this role Have the instances bootstrap get the certificate from Amazon 53 upon first boot.
B. Embed a certificate into the Amazon Machine Image that is used by the Auto Scaling group Have the launched instances generate a certificate signature request with the instance's assigned instance- id to the Key management service for signature.
C. Configure the Auto Scaling group to send an SNS notification of the launch of a new instance to the trusted key management servic
D. Have the Key management service generate a signed certificate and send it directly to the newly launched instance.
E. Configure the launched instances to generate a new certificate upon first boot Have the Key management service poll the AutoScaling group for associated instances and send new instances acertificate signature (hat contains the specific instance-i

Answer: A

NEW QUESTION 200

Your company has recently extended its datacenter into a VPC on AVVS to add burst computing capacity as needed Members of your Network Operations Center need to be able to go to the AWS Management Console and administer Amazon EC2 instances as necessary You don't want to create new IAM users for each NOC member and make those users sign in again to the AWS Management Console Which option below will meet the needs for your NOC members?

A. Use OAuth 2 0 to retrieve temporary AWS security credentials to enable your NOC members to sign in to the AVVS Management Console.
B. Use web Identity Federation to retrieve AWS temporary security credentials to enable your NOC members to sign in to the AWS Management Console.
C. Use your on-premises SAML 2.0-compliant identity provider (IOP) to grant the NOC members federated access to the AWS Management Console via the AWS sing Ie sign-on (550) endpoint.
D. Use your on-premises SAML2.0-comp|iam identity provider (IOP) to retrieve temporary security credentials to enable NOC members to sign in to the AWS Management Console.

Answer: D

NEW QUESTION 204
You are designing an SSUTLS solution that requires HTIPS clients to be authenticated by the Web server using client certificate authentication. The solution must be resilient.
Which of the following options would you consider for configuring the web server infrastructure? (Choose 2 answers)

A. Configure ELB with TCP listeners on TCP/4d3. And place the Web servers behind it.
B. Configure your Web servers with EIPS Place the Web servers in a Route53 Record Set and configure health checks against all Web servers.
C. Configure ELB with HTIPS listeners, and place the Web servers behind it.
D. Configure your web servers as the origins for a Cloud Front distributio
E. Use custom SSL certificates on your Cloud Front distribution.

Answer: AB

NEW QUESTION 207
You require the ability to analyze a customer's clickstream data on a website so they can do behavioral analysis. Your customer needs to know what sequence of pages and ads their customer clicked on. This data will be used in real time to modify the page layouts as customers click through the site to increase stickiness and advertising click-through. Which option meets the requirements for captioning and analyzing this data?

A. Log clicks in weblogs by URL store to Amazon 53, and then analyze with Elastic MapReduce
B. Push web clicks by session to Amazon Kinesis and analyze behavior using Kinesis workers
C. Write click events directly to Amazon Redshift and then analyze with SQL
D. Publish web clicks by session to an Amazon SQS queue men periodically drain these events to Amazon RDS and analyze with sol

Answer: B

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

NEW QUESTION 212
An AWS customer runs a public blogging website. The site users upload two million blog entries a month. The average blog entry size is 200 KB. The access rate to blog entries drops to negligible 6 months after publication and users rarely access a blog entry 1 year after publication. Additionally, blog entries have a high update rate during the first 3 months following publication, this drops to no updates after 6 months. The customer wants to use CIoudFront to improve his user's load times.
Which of the following recommendations would you make to the customer?

A. Duplicate entries into two different buckets and create two separate CIoudFront distributions where 53 access is restricted only to Cloud Front identity
B. Create a CIoudFront distribution with "US" Europe price class for US/ Europe users and a different CIoudFront distribution with AI I Edge Locations' for the remaining users.
C. Create a CIoudFront distribution with 53 access restricted only to the CIoudFront identity and partition the blog entry's location in 53 according to the month it was uploaded to be used with CIoudFront behaviors.
D. Create a CIoudFronI distribution with Restrict Viewer Access Forward Query string set to true and minimum TTL of 0.

Answer: C

NEW QUESTION 213
Your company is getting ready to do a major public announcement of a social media site on AWS. The website is running on EC2 instances deployed across multiple Availability Zones with a MuIti-AZ RDS MySQL Extra Large DB Instance. The site performs a high number of small reads and writes per second and relies on an eventual consistency model. After comprehensive tests you discover that there is read contention on RDS MySQL. Which are the best approaches to meet these requirements? (Choose 2 answers)

A. Deploy EIasticCache in-memory cache running in each availability zone
B. Implement sharding to distribute load to multiple RDS MySQL instances
C. Increase the RDS MySQL Instance size and Implement provisioned IQPS
D. Add an RDS MySQL read replica in each availability zone

Answer: AC

NEW QUESTION 214
A read only news reporting site with a combined web and application tier and a database tier that receives large and unpredictable traffic demands must be able to respond to these traffic fluctuations automatically. What AWS services should be used meet these requirements?

A. Stateless instances for the web and application tier synchronized using Elasticache Memcached in an autoscaimg group monitored with CIoudWatc
B. And RDSwith read replicas.
C. Stateful instances for the web and application tier in an autoscaling group monitored with CIoudWatch and RDS with read replicas.
D. Stateful instances for the web and application tier in an autoscaling group monitored with CIoudWatc

E. And multi-AZ RDS.
F. Stateless instances for the web and application tier synchronized using EIastiCache Memcached in an autoscaling group monitored with CIoudWatch and multi-AZ RDS.

Answer: A

NEW QUESTION 217
Your customer is willing to consolidate their log streams (access logs application logs security logs etc.) in one single system. Once consolidated, the customer wants to analyze these logs in real time based on heuristics. From time to time, the customer needs to validate heuristics, which requires going back to data samples extracted from the last 12 hours?
What is the best approach to meet your customer's requirements?

A. Send all the log events to Amazon SQ
B. Setup an Auto Scaling group of EC2 sewers to consume the logs and apply the heuristics.
C. Send all the log events to Amazon Kinesis develop a client process to apply heuristics on the logs
D. Configure Amazon Cloud Trail to receive custom logs, use EMR to apply heuristics the logs
E. Setup an Auto Scaling group of EC2 syslogd servers, store the logs on 53 use EMR to apply heuristics on the logs

Answer: B

Explanation:
The throughput of an Amazon Kinesis stream is designed to scale without limits via increasing the number of shards within a stream. However, there are certain limits you should keep in mind while using Amazon Kinesis Streams:
By default, Records of a stream are accessible for up to 24 hours from the time they are added to the stream. You can raise this limit to up to 7 days by enabling extended data retention.
The maximum size of a data blob (the data payload before Base64-encoding) within one record is 1 megabyte (MB).
Each shard can support up to 1000 PUT records per second.
For more information about other API level limits, see Amazon Kinesis Streams Limits.

NEW QUESTION 222
You deployed your company website using Elastic Beanstalk and you enabled log file rotation to 53. An Elastic Map Reduce job is periodically analyzing the logs on 53 to build a usage dashboard that you share with your CIO.
You recently improved overall performance of the website using Cloud Front for dynamic content delivery and your website as the origin.
After this architectural change, the usage dashboard shows that the traffic on your website dropped by an order of magnitude. How do you fix your usage dashboard'?

A. Enable Cloud Front to deliver access logs to 53 and use them as input of the Elastic Map Reduce job.
B. Turn on Cloud Trail and use trail log tiles on 53 as input of the Elastic Map Reduce job
C. Change your log collection process to use Cloud Watch ELB metrics as input of the Elastic MapReducejob
D. Use Elastic Beanstalk "Rebuild Environment" option to update log delivery to the Elastic Map Reduce job.
E. Use Elastic Beanstalk 'Restart App server(s)" option to update log delivery to the Elastic Map Reduce job.

Answer: D

NEW QUESTION 223
You are running a successful multitier web application on AWS and your marketing department has asked you to add a reporting tier to the application. The reporting tier will aggregate and publish status reports every 30 minutes from user-generated information that is being stored in your web application s database. You are currently running a Multi-AZ RDS MySQL instance for the database tier. You also have implemented Elasticache as a database caching layer between the application tier and database tier. Please select the answer that will allow you to successful ly implement the reporting tier with as little impact as possible to your database.

A. Continually send transaction logs from your master database to an 53 bucket and generate the reports off the 53 bucket using 53 byte range request s.
B. Generate the reports by querying the synchronously replicated standby RDS MySQL instance maintained through Multi-AZ.
C. Launch a RDS Read Replica connected to your Multi AZ master database and generate reports by querying the Read Replica.
D. Generate the reports by querying the EIastiCache database caching tie

Answer: C

Explanation:
Amazon RDS allows you to use read replicas with MuIti-AZ deployments. In Multi-AZ deployments for MySQL, Oracle, SQL Server, and PostgreSQL, the data in your primary DB Instance is synchronously replicated to a standby instance in a different Availability Zone (AZ). Because of their synchronous replication, MuIti-AZ deployments for these engines offer greater data durability benefits than do read replicas. (In all Amazon RDS for Aurora deployments, your data is automatically replicated across 3 Availability Zones.)
You can use MuIti-AZ deployments and read replicas in conjunction to enjoy the complementary benefits of each. You can simply specify that a given Multi-AZ deployment is the source DB Instance for your Read replicas. That way you gain both the data durability and availability benefits of Multi -AZ deployments and the read scaling benefits of read replicas.
Note that for MuIti-AZ deployments, you have the option to create your read replica in an AZ other than that of the primary and the standby for even more redundancy. You can identify the AZ corresponding to your standby by looking at the "Secondary Zone" field of your DB Instance in the AWS Management Console.

NEW QUESTION 224
Fill in the blanks: Resources that are created in AWS are identified by a unique identifier called an

A. Amazon Resource Number
B. Amazon Resource Nametag
C. Amazon Resource Name
D. Amazon Resource Namespace

Answer: C

NEW QUESTION 228
While creating an Amazon RDS DB, your first task is to set up a DB _ that controls what IP addresses or EC2 instances have access to your DB Instance.

A. Security Pool
B. Secure Zone
C. Security Token Pool
D. Security Group

Answer: D

NEW QUESTION 229
By default, EBS volumes that are created and attached t o an instance at launch are deleted when t hat instance is terminated. You can modify this behavior by changing the value of the flag _ to false when you launch the instance

A. Delete On Termination
B. Remove On Deletion
C. Remove On Termination
D. Terminate On Deletion

Answer: A

NEW QUESTION 234
True or False: When using IAM to control access to your RDS resources, the key names that can be used are case sensitive. For example, aws:CurrentTime is NOT equivalent to AWS:currenttime.

A. TRUE
B. FALSE

Answer: A

NEW QUESTION 236
What will be the status of the snapshot until the snapshot is complete.

A. running
B. working
C. progressing
D. pending

Answer: D

NEW QUESTION 241
True or False: Automated backups are enabled by default for a new DB Instance.

A. TRUE
B. FALSE

Answer: A

NEW QUESTION 245
What does the AWS Storage Gateway provide?

A. It allows to integrate on-premises IT environments with Cloud Storage.
B. A direct encrypted connection to Amazon 53.
C. It's a backup solution that provides an on-premises Cloud storage.
D. It provides an encrypted SSL endpoint for backups in the Clou

Answer: A

NEW QUESTION 246
Amazon RDS automated backups and DB Snapshots are currently supported for only the _ _ storage engine

A. InnoDB
B. MyISAM

Answer: A

NEW QUESTION 249
What are the two permission types used by AWS'?

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

Answer: D

NEW QUESTION 250
In the Amazon cloudwatch, which metric should I be checking to ensure that your DB Instance has enough free storage space?

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

Answer: B

NEW QUESTION 252
Amazon RDS DB snapshots and automated backups are stored in

A. Amazon 53
B. Amazon ECS Volume
C. Amazon RDS
D. Amazon EMR

Answer: A

NEW QUESTION 256
What is the maximum key length of a tag'?

A. 512 Unicode characters
B. 64 Unicode characters
C. 256 Unicode characters
D. 128 Unicode characters

Answer: D

NEW QUESTION 259
Changes to the backup window take effect _ _

A. from the next billing cycle
B. after 30 minutes
C. immediately
D. after 24 hours

Answer: C

NEW QUESTION 264
What are the two types of licensing options available for using Amazon RDS for Oracle?

A. BYOL and Enterprise License
B. BYOL and License Included
C. Enterprise License and License Included
D. Role based License and License Included

Answer: B

NEW QUESTION 267
What is the durability of 53 RRS?

A. 99.99%
B. 99.95%
C. 99.995%
D. 99.999999999%

Answer: A

NEW QUESTION 270
Is Federated Storage Engine currently supported by Amazon RDS for MySQL?

A. Only for Oracle RDS instances
B. No
C. Yes
D. Only in VPC

Answer: B

NEW QUESTION 271
While launching an RDS DB instance, on which page I can select the Availability Zone?

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

Answer: D

NEW QUESTION 275
To view information about an Amazon EBS volume, open the Amazon EC2 console at https://console.aws.amazon.com/ec2/, click in the Navigation pane.

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

Answer: D

NEW QUESTION 277
Read Replicas require a transactional storage engine and are only supported for the _ _ storage engine

A. OracleISAM
B. MSSQLDB
C. InnoDB
D. IV|y|SAIV|

Answer: C

NEW QUESTION 280
When running my DB Instance as a Multi-AZ deployment, can I use the standby for read or write operations?

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

Answer: D

NEW QUESTION 282
Are Resenred Instances available for Multi-AZ Deployments?

A. Only for Cluster Compute instances
B. Yes for all instance types
C. Only for M3 instance types
D. No

Answer: B

NEW QUESTION 284
What is the command line instruction for running the remote desktop client in Windows?

A. desk.cpl
B. mstsc

Answer: B

NEW QUESTION 287
MySQL installations default to port _.

A. A.3306B.443
B. 80
C. 1158

Answer: A

NEW QUESTION 288
What does Amazon Cloud Formation provide?

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

Answer: D

NEW QUESTION 290
What can I access by visiting the URL: http:/ /status.aws.amazon.com/?

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

Answer: C

NEW QUESTION 295
What is the type of monitoring data (for Amazon EBS volumes) which is available automatically in 5- minute periods at no charge called?

A. Basic
B. Primary
C. Detailed
D. Local

Answer: A

NEW QUESTION 299
When you use the AWS Management Console to delete an IAM user, IAM also deletes any signing certificates and any access keys belonging to the user.

A. FALSE
B. This is configurable
C. TRUE

Answer: C

NEW QUESTION 303
When automatic failover occurs, Amazon RDS will emit a DB Instance event to inform you that automatic failover occurred. You can use the to ret urn information about events related to your DB Instance

A. FetchFai|ure
B. DescriveFai|ure
C. DescribeEvents
D. FetchEvents

Answer: C

NEW QUESTION 304
Select the correct statement:

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

Answer: C

NEW QUESTION 306
Is the encryption of connections between my application and my DB Instance using SSL for the MySQL server engines available?

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

Answer: A

NEW QUESTION 310
What does Amazon Elastic Beanstalk provide?

A. An application container on top of Amazon Web Services.
B. A scalable storage appliance on top of Amazon Web Services.
C. A scalable cluster of EC2 instances.
D. A service by this name doesn't exis

Answer: C

NEW QUESTION 313
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

AWS-Solution-Architect-Associate Practice Exam Features:

* AWS-Solution-Architect-Associate Questions and Answers Updated Frequently

* AWS-Solution-Architect-Associate Practice Questions Verified by Expert Senior Certified Staff

* AWS-Solution-Architect-Associate Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* AWS-Solution-Architect-Associate Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year