AWS-Solution-Architect-Associate Dumps

Amazon AWS Certified Solutions Architect - Associate

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

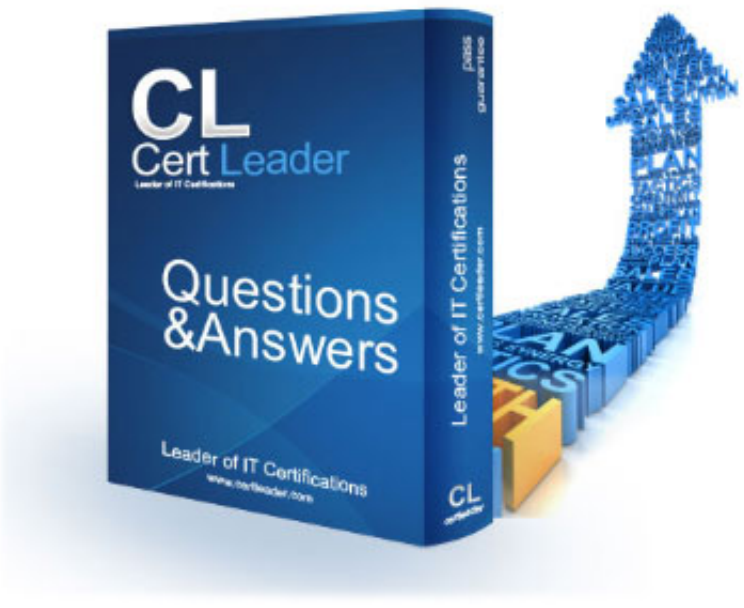

NEW QUESTION 1
You are trying to launch an EC2 instance, however the instance seems to go into a terminated status immediately. What would probably not be a reason that this is happening?

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

Answer: C

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

NEW QUESTION 2
Amazon EBS provides the ability to create backups of any Amazon EC2 volume into what is known as

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

Answer: A

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

NEW QUESTION 3
One of the criteria for a new deployment is that the customer wants to use AWS Storage Gateway. However you are not sure whether you should use gateway-cached volumes or gateway-stored volumes or even what the differences are. Which statement below best describes those differences?

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

Answer: A

Explanation:
Volume gateways provide cloud-backed storage volumes that you can mount as Internet Small Computer System Interface (iSCSI) devices from your on-premises application sewers. The gateway supports the following volume configurations:
Gateway-cached volumes — You store your data in Amazon Simple Storage Service (Amazon S3) and retain a copy of frequently accessed data subsets locally. Gateway-cached volumes offer a substantial cost savings on primary storage and minimize the need to scale your storage on-premises. You also retain low-latency access to your frequently accessed data.
Gateway-stored volumes — If you need low-latency access to your entire data set, you can configure your on-premises gateway to store all your data locally and then asynchronously back up point-in-time snapshots of this data to Amazon S3. This configuration provides durable and inexpensive off-site backups that you can recover to your local data center or Amazon EC2. For example, if you need replacement capacity for disaster recovery, you can recover the backups to Amazon EC2.
Reference: http://docs.aws.amazon.com/storagegateway/latest/userguide/volume-gateway.html

NEW QUESTION 4
An edge location refers to which Amazon Web Service?

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

Answer: C

Explanation:
Amazon CIoudFront is a content distribution network. A content delivery network or content distribution network (CDN) is a large distributed system of sewers deployed in multiple data centers across the world. The location of the data center used for CDN is called edge location.
Amazon CIoudFront can cache static content at each edge location. This means that your popular static content (e.g., your site's logo, navigational images, cascading style sheets, JavaScript code, etc.) will be available at a nearby edge location for the browsers to download with low latency and improved performance for viewers. Caching popular static content with Amazon CIoudFront also helps you offload requests for such files from your origin sever — CIoudFront serves the cached copy when available and only makes a request to your origin server if the edge location receMng the browser's request does not have a copy of the file.
Reference: http://aws.amazon.com/c|oudfront/

NEW QUESTION 5
You are looking at ways to improve some existing infrastructure as it seems a lot of engineering resources are being taken up with basic management and monitoring tasks and the costs seem to be excessive.
You are thinking of deploying Amazon E|asticCache to help. Which of the following statements is true in regards to EIasticCache?

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

Answer: D

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

NEW QUESTION 6
Your supervisor has asked you to build a simple file synchronization service for your department. He doesn't want to spend too much money and he wants to be notified of any changes to files by email. What do you think would be the best Amazon service to use for the email solution?

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

Answer: A

Explanation:
File change notifications can be sent via email to users following the resource with Amazon Simple Email Service (Amazon SES), an easy-to-use, cost-effective email solution.
Reference: http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_fiIesync_08.pdf

NEW QUESTION 7
Amazon EC2 provides a . It is an HTTP or HTTPS request that uses the HTTP verbs GET or POST.

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

Answer: C

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

NEW QUESTION 8
In Amazon AWS, which of the following statements is true of key pairs?

A. Key pairs are used only for Amazon SDKs.
B. Key pairs are used only for Amazon EC2 and Amazon CIoudFront.
C. Key pairs are used only for Elastic Load Balancing and AWS IAM.
D. Key pairs are used for all Amazon service

Answer: B

Explanation:
Key pairs consist of a public and private key, where you use the private key to create a digital signature, and then AWS uses the corresponding public key to validate the signature. Key pairs are used only for Amazon EC2 and Amazon CIoudFront.
Reference: http://docs.aws.amazon.com/generaI/latest/gr/aws-sec-cred-types.html

NEW QUESTION 9
You need to migrate a large amount of data into the cloud that you have stored on a hard disk and you decide that the best way to accomplish this is with AWS Import/Export and you mail the hard disk to AWS. Which of the following statements is incorrect in regards to AWS Import/Export?

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

Answer: C

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

NEW QUESTION 10
Which of the following is true of Amazon EC2 security group?

A. You can modify the outbound rules for EC2-Classic.
B. You can modify the rules for a security group only if the security group controls the traffic for just one instance.
C. You can modify the rules for a security group only when a new instance is created.
D. You can modify the rules for a security group at any tim

Answer: D

Explanation:
A security group acts as a virtual firewall that controls the traffic for one or more instances. When you launch an instance, you associate one or more security groups with the instance. You add rules to each security group that allow traffic to or from its associated instances. You can modify the rules for a security group at any time; the new rules are automatically applied to all instances that are associated with the security group.
Reference: http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/using-network-security.html

NEW QUESTION 10
You have been asked to build a database warehouse using Amazon Redshift. You know a little about it, including that it is a SQL data warehouse solution, and uses industry standard ODBC and JDBC connections and PostgreSQL drivers. However you are not sure about what sort of storage it uses for database tables. What sort of storage does Amazon Redshift use for database tables?

A. InnoDB Tables
B. NDB data storage
C. Columnar data storage
D. NDB CLUSTER Storage

Answer: C

Explanation:
Amazon Redshift achieves efficient storage and optimum query performance through a combination of massively parallel processing, columnar data storage, and very efficient, targeted data compression encoding schemes.
Columnar storage for database tables is an important factor in optimizing analytic query performance because it drastically reduces the overall disk I/O requirements and reduces the amount of data you need to load from disk.
Reference: http://docs.aws.amazon.com/redshift/latest/dg/c_co|umnar_storage_disk_mem_mgmnt.html

NEW QUESTION 12
You are building infrastructure for a data warehousing solution and an extra request has come through that there will be a lot of business reporting queries running all the time and you are not sure if your current DB instance will be able to handle it. What would be the best solution for this?

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

Answer: B

Explanation:
Read Replicas make it easy to take advantage of MySQL’s built-in replication functionality to elastically scale out beyond the capacity constraints of a single DB Instance for read-heavy database workloads. There are a variety of scenarios where deploying one or more Read Replicas for a given source DB Instance may make sense. Common reasons for deploying a Read Replica include:
Scaling beyond the compute or I/O capacity of a single DB Instance for read-heavy database workloads. This excess read traffic can be directed to one or more Read Replicas.
Serving read traffic while the source DB Instance is unavailable. If your source DB Instance cannot take I/O requests (e.g. due to I/O suspension for backups or scheduled maintenance), you can direct read traffic to your Read Replica(s). For this use case, keep in mind that the data on the Read Replica may be "staIe" since the source DB Instance is unavailable.
Business reporting or data warehousing scenarios; you may want business reporting queries to run against a Read Replica, rather than your primary, production DB Instance.
Reference: https://aws.amazon.com/rds/faqs/

NEW QUESTION 17
In DynamoDB, could you use IAM to grant access to Amazon DynamoDB resources and API actions?

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

Answer: D

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

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

NEW QUESTION 18
An online gaming site asked you if you can deploy a database that is a fast, highly scalable NoSQL database service in AWS for a new site that he wants to build. Which database should you recommend?

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

Answer: A

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

NEW QUESTION 21
You need to change some settings on Amazon Relational Database Service but you do not want the database to reboot immediately which you know might happen depending on the setting that you change. Which of the following will cause an immediate DB instance reboot to occur?

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

Answer: A

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

NEW QUESTION 24
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 27
What would be the best way to retrieve the public IP address of your EC2 instance using the CLI?

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

Answer: D

Explanation:
To determine your instance's public IP address from within the instance, you can use instance metadata. Use the following command to access the public IP address: For Linux use, $ curl
http://169.254.169.254/latest/meta-data/public-ipv4, and for Windows use, $ wget http://169.254.169.254/latest/meta-data/public-ipv4.

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

NEW QUESTION 28
Having set up a website to automatically be redirected to a backup website if it fails, you realize that there are different types of failovers that are possible. You need all your resources to be available the majority of the time. Using Amazon Route 53 which configuration would best suit this requirement?

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

Answer: A

Explanation:
You can set up a variety of failover configurations using Amazon Route 53 alias: weighted, latency, geolocation routing, and failover resource record sets.
Active-active failover: Use this failover configuration when you want all of your resources to be available the majority of the time. When a resource becomes unavailable, Amazon Route 53 can detect that it's unhealthy and stop including it when responding to queries.
Active-passive failover: Use this failover configuration when you want a primary group of resources to be available the majority of the time and you want a secondary group of resources to be on standby in case all of the primary resources become unavailable. When responding to queries, Amazon Route 53 includes only the healthy primary resources. If all of the primary resources are unhealthy, Amazon Route 53 begins to include only the healthy secondary resources in response to DNS queries.
Active-active-passive and other mixed configurations: You can combine alias and non-alias resource record sets to produce a variety of Amazon Route 53 behaviors.
Reference: http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover.html

NEW QUESTION 30
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 35
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 38
A user is accessing an EC2 instance on the SSH port for IP 10.20.30.40. Which one is a secure way to configure that the instance can be accessed only from this IP?

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

Answer: B

Explanation:
In AWS EC2, while configuring a security group, the user needs to specify the IP address in CIDR notation. The CIDR IP range 10.20.30.40/32 says it is for a single IP 10.20.30.40. If the user specifies the IP as 10.20.30.40 only, the security group will not accept and ask it in a CIRD format.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html

NEW QUESTION 39

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

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

Answer: A

Explanation:
The user can join multiple provisioned IOPS volumes together in a RAID 1 configuration to achieve better fault tolerance. RAID 1 does not provide a write performance improvement; it requires more bandwidth than non-RAID configurations since the data is written simultaneously to multiple volumes.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/raid-config.html

NEW QUESTION 40
A user has created a subnet in VPC and launched an EC2 instance within it. The user has not selected the option to assign the IP address while launching the instance. The user has 3 elastic IPs and is trying to assign one of the Elastic IPs to the VPC instance from the console. The console does not show any instance in the IP assignment screen. What is a possible reason that the instance is unavailable in the assigned IP console?

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

Answer: D

Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. When the user is launching an instance he needs to select an option which attaches a public IP to the instance. If the user has not selected the option to attach the public IP then it will only have a private IP when launched. If the user wants to connect to
an instance from the internet he should create an elastic IP with VPC. If the elastic IP is a part of EC2
Classic it cannot be assigned to a VPC instance.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/GettingStartedGuide/LaunchInstance.html

NEW QUESTION 45
A user has launched one EC2 instance in the US East region and one in the US West region. The user has launched an RDS instance in the US East region. How can the user configure access from both the EC2 instances to RDS?

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

Answer: C

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

NEW QUESTION 49
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 52
The common use cases for DynamoDB Fine-Grained Access Control (FGAC) are cases in which the end user wants .

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

Answer: D

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

NEW QUESTION 54
A user comes to you and wants access to Amazon CIoudWatch but only wants to monitor a specific LoadBaIancer. Is it possible to give him access to a specific set of instances or a specific LoadBaIancer?

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

Answer: A

Explanation:
Amazon CIoudWatch integrates with AWS Identity and Access Management (IAM) so that you can specify which CIoudWatch 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 CIoudWatch data for specific resources. For example, you can't give a user access to CIoudWatch data for only a specific set of instances or a specific LoadBaIancer. Permissions granted using IAM cover all the cloud resources you use with CIoudWatch. In addition, you can't use IAM roles with the Amazon CIoudWatch command line tools.
Using Amazon CIoudWatch with IAM doesn't change how you use CIoudWatch. There are no changes to CIoudWatch actions, and no new CIoudWatch actions related to users and access control.
Reference: http://docs.aws.amazon.com/AmazonC|oudWatch/latest/DeveloperGuide/UsingIAM.htmI

NEW QUESTION 55
A user is planning to make a mobile game which can be played online or offline and will be hosted on EC2.
The user wants to ensure that if someone breaks the highest score or they achieve some milestone they can inform all their colleagues through email. Which of the below mentioned AWS services helps achieve this goal?

A. AWS Simple Workflow Service.
B. AWS Simple Email Service.
C. Amazon Cognito
D. AWS Simple Queue Servic

Answer: B

Explanation:
Amazon Simple Email Service (Amazon SES) is a highly scalable and cost-effective email-sending service for businesses and developers. It integrates with other AWS services, making it easy to send emails from applications that are hosted on AWS.
Reference: http://aws.amazon.com/ses/faqs/

NEW QUESTION 59
You have multiple VPN connections and want to provide secure communication between sites using the AWS VPN CIoudHub. Which statement is the most accurate in describing what you must do to set this up correctly?

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

Answer: A

Explanation:
If you have multiple VPN connections, you can provide secure communication between sites using the AWS VPN 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'd like to implement a convenient, potentially low-cost hub-and-spoke model for primary or backup connectMty between these remote offices.
To use the AWS VPN CIoudHub, you must create a virtual private gateway with multiple customer gateways, each with unique Border Gateway Protocol (BGP) Autonomous System Numbers (ASNs). Customer gateways advertise the appropriate routes (BGP prefixes) over their VPN connections. These routing advertisements are received and re-advertised to each BGP peer, enabling each site to send data to and receive data from the other sites. The routes for each spoke must have unique ASNs and the sites must not have overlapping IP ranges. Each site can also send and receive data from the VPC as if they were using a standard VPN connection.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPN_CIoudHub.htmI

NEW QUESTION 61
You need to create an Amazon Machine Image (AM) for a customer for an application which does not appear to be part of the standard AWS AM template that you can see in the AWS console. What are the alternative possibilities for creating an AM on AWS?

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

Answer: B

Explanation:
You can purchase an AMIs from a third party, including AMIs that come with service contracts from organizations such as Red Hat. You can also create an AMI and sell it to other Amazon EC2 users. After you create an AMI, you can keep it private so that only you can use it, or you can share it with a specified list of AWS accounts. You can also make your custom AMI public so that the community can
use it. Building a safe, secure, usable AMI for public consumption is a fairly straightforward process, if you follow a few simple guidelines.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AMIs.htm|

NEW QUESTION 66
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 71
You have been asked to tighten up the password policies in your organization after a serious security breach, so you need to consider every possible security measure. Which of the following is not an account password policy for IAM Users that can be set?

A. Force IAM users to contact an account administrator when the user has allowed his or her password to expue.
B. A minimum password length.
C. Force IAM users to contact an account administrator when the user has entered his password incorrectly.
D. Prevent IAM users from reusing previous password

Answer: C

Explanation:
IAM users need passwords in order to access the AWS Management Console. (They do not need passwords if they will access AWS resources programmatically by using the CLI, AWS SDKs, or the APIs.)
You can use a password policy to do these things: Set a minimum password length.
Require specific character types, including uppercase letters, lowercase letters, numbers, and non-alphanumeric characters. Be sure to remind your users that passwords are case sensitive. Allow all IAM users to change their own passwords.
Require IAM users to change their password after a specified period of time (enable password expiration). Prevent IAM users from reusing previous passwords.
Force IAM users to contact an account administrator when the user has allowed his or her password to expue.
Reference: http://docs.aws.amazon.com/|AM/Iatest/UserGuide/Using_ManagingPasswordPoIicies.htm|

NEW QUESTION 72
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 75
You have just finshed setting up an advertisement server in which one of the obvious choices for a service was Amazon Elastic Map Reduce(EMR) and are now troubleshooting some weird cluster states that you are seeing. Which of the below is not an Amazon EMR cluster state?

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

Answer: B

Explanation:

Amazon Elastic Map Reduce (EMR) is a web service that enables businesses, researchers, data analysts, and developers to easily and cost-effectively process vast amounts of data.
Amazon EMR historically referred to an Amazon EMR cluster (and all processing steps assigned to it) as a "c|uster". Every cluster has a unique identifier that starts with "j-".
The different cluster states of an Amazon EMR cluster are listed below. STARTING — The cluster provisions, starts, and configures EC2 instances.
BOOTSTRAPPING — Bootstrap actions are being executed on the cluster. RUNNING — A step for the cluster is currently being run.
WAITING — The cluster is currently active, but has no steps to run. TERMINATING - The cluster is in the process of shutting down. TERMINATED - The cluster was shut down without error. TERMINATED_W|TH_ERRORS - The cluster was shut down with errors.
Reference: https://aws.amazon.com/elasticmapreduce/faqs/

NEW QUESTION 80
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 CIoudTraiI 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 85
Which of the following would you use to list your AWS Import/Exportjobs?

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

Answer: C

Explanation:
You can list AWS Import/Export jobs with the ListJobs command using the command line client or REST API.
Reference: http://docs.aws.amazon.com/AWSImportExport/latest/DG/ListingYourJobs.html

NEW QUESTION 87
Mike is appointed as Cloud Consultant in Netcrak Inc. Netcrak has the following VPCs set-up in the US East Region:
A VPC with CIDR block 10.10.0.0/16, a subnet in that VPC with CIDR block 10.10.1.0/24 A VPC with CIDR block 10.40.0.0/16, a subnet in that VPC with CIDR block 10.40.1.0/24
Netcrak Inc is trying to establish network connection between two subnets, a subnet with CIDR block 10.10.1.0/24 and another subnet with CIDR block 10.40.1.0/24. Which one of the following solutions should Mke recommend to Netcrak Inc?

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

Answer: C

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

NEW QUESTION 90
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 94
Any person or application that interacts with AWS requires security credentials. AWS uses these credentials to identify who is making the call and whether to allow the requested access. You have just set up a VPC network for a client and you are now thinking about the best way to secure this network. You set up a security group called vpcsecuritygroup. Which following statement is true in respect to the initial settings that will be applied to this security group if you choose to use the default settings for this group?

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

Answer: B

Explanation:
Amazon VPC provides advanced security features such as security groups and network access control lists to enable inbound and outbound filtering at the instance level and subnet level.
AWS assigns each security group a unique ID in the form sg-xxxxxxxx. The following are the initial settings for a security group that you create:
Allow no inbound traffic Allow all outbound traffic
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html

NEW QUESTION 95
Having just set up your first Amazon Virtual Private Cloud (Amazon VPC) network, which defined a default network interface, you decide that you need to create and attach an additional network interface, known as an elastic network interface (ENI) to one of your instances. Which of the following statements is true regarding attaching network interfaces to your instances in your VPC?

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

Answer: C

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

NEW QUESTION 100
An organization has a statutory requirement to protect the data at rest for the S3 objects. Which of the below mentioned options need not be enabled by the organization to achieve data security?

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

Answer: D

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

NEW QUESTION 101
In Amazon CIoudFront, if you use Amazon EC2 instances and other custom origins with CIoudFront, 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 CIoudWatch metrics
D. specify the URL of the load balancer for the domain name of your origin server

Answer: D

Explanation:
In Amazon CIoudFront, 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 103
Which of the following features are provided by Amazon EC2?

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

Answer: B

Explanation:
Amazon EC2 provides the following features:
· Virtual computing environments, known as instances;
· Pre-configured templates for your instances, known as Amazon Nlachine Images (AMIs), that package the bits you need for your server (including the operating system and additional software)
· Various configurations of CPU, memory, storage, and networking capacity for your instances, known as instance types
· Secure login information for your instances using key pairs (AWS stores the public key, and you store the private key in a secure place)
· Storage volumes for temporary data that's deleted when you stop or terminate your instance, known as instance store volumes
· Persistent storage volumes for your data using Amazon Elastic Block Store (Amazon EBS), known as Amazon EBS volumes
· Mulliple physical locations for your resources, such as instances and Amazon EBS volumes, known as regions and Availability Zones
· A firewall that enables you to specify the protocols, ports, and source IP ranges that can reach your instances using security groups
· Static IP addresses for dynamic cloud computing, known as Elastic IP addresses
· Metadata, known as tags, that you can create and assign to your Amazon EC2 resources
· Virtual networks you can create that are logically isolated from the rest of the AWS cloud, and that you can optionally connect to your own network, known as virtual private clouds (VPCs).
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/concepts.html

NEW QUESTION 104
A user is sending bulk emails using AWS SES. The emails are not reaching some of the targeted audience because they are not authorized by the ISPs. How can the user ensure that the emails are all delivered?

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

Answer: A

Explanation:
Domain Keys Identified Mail (DKIM) is a standard that allows senders to sign their email messages and ISPs, and use those signatures to verify that those messages are legitimate and have not been modified by a third party in transit.
Reference: http://docs.aws.amazon.com/ses/latest/DeveloperGuide/dkim.html

NEW QUESTION 109
In AWS CIoudHSM, 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 CIoudHSM, 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 110
You are architecting a highly-scalable and reliable web application which will have a huge amount of content .You have decided to use Cloudfront as you know it will speed up distribution of your static and dynamic web content and know that Amazon C|oudFront integrates with Amazon CIoudWatch metrics so that you can monitor your web application. Because you live in Sydney you have chosen the the Asia Pacific (Sydney) region in the AWS console. However you have set up this up but no CIoudFront metrics seem to be appearing in the CIoudWatch console. What is the most likely reason from the possible choices below for this?

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

Answer: C

Explanation:
CIoudFront is a global service, and metrics are available only when you choose the US East (N. Virginia) region in the AWS console. If you choose another region, no CIoudFront metrics will appear in the CIoudWatch console.
Reference:
http://docs.aws.amazon.com/AmazonCIoudFront/latest/Deve|operGuide/monitoring-using-cloudwatch.ht ml

NEW QUESTION 113
After a major security breach your manager has requested a report of all users and their credentials in AWS. You discover that in IAM you can generate and download a credential report that lists all users in your account and the status of their various credentials, including passwords, access keys, MFA devices, and signing certificates. Which following statement is incorrect in regards to the use of credential reports?

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

Answer: A

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

NEW QUESTION 116
A user is planning a highly available application deployment with EC2. Which of the below mentioned options will not help to achieve HA?

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

Answer: B

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

NEW QUESTION 119
After deploying a new website for a client on AWS, he asks if you can set it up so that if it fails it can be automatically redirected to a backup website that he has stored on a dedicated server elsewhere. You are wondering whether Amazon Route 53 can do this. Which statement below is correct in regards to Amazon Route 53?

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

Answer: B

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

NEW QUESTION 124
You need to create a management network using network interfaces for a virtual private cloud (VPC) network. Which of the following statements is incorrect pertaining to Best Practices for Configuring Network Interfaces.

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

Answer: D

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

NEW QUESTION 128
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
Can you encrypt EBS volumes?

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

Answer: A

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

NEW QUESTION 138
A user is running a webserver on EC2. The user wants to receive the SMS when the EC2 instance utilization is above the threshold limit. Which AWS services should the user configure in this case?

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

Answer: B

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

NEW QUESTION 140
Your manager has come to you saying that he is very confused about the bills he is receMng from AWS as he is getting different bills for every user and needs you to look into making it more understandable. Which of the following would be the best solution to meet his request?

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

Answer: B

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

NEW QUESTION 144
When controlling access to Amazon EC2 resources, each Amazon EBS Snapshot has a attribute that controls which AWS accounts can use the snapshot.

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

Answer:

A

Explanation:
Each Amazon EBS Snapshot has a 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 149
A 3-tier e-commerce web application is current deployed on-premises and will be migrated to AWS for
greater scalability and elasticity The web server currently shares read-only data using a network distributed file system The app server tier uses a clustering mechanism for discovery and shared session state that depends on I P multicast The database tier uses shared-storage clustering to provide database fail over capability, and uses several read slaves for scaling Data on all sewers and the distributed file system directory is backed up weekly to off-site tapes
Which AWS storage and database architecture meets the requirements of the application?

A. Web sewers: store read-only data in 53, and copy from 53 to root volume at boot tim
B. App servers: share state using a combination of DynamoDB and IP unicas
C. Database: use RDS with multi-AZ deployment and one or more read replica
D. Backup: web servers, app servers, and database backed up weekly to Glacier using snapshots.
E. Web sewers: store read-only data in an EC2 NFS server, mount to each web server at boot tim
F. App servers: share state using a combination of DynamoDB and IP multicas
G. Database: use RDS with multi-AZ deployment and one or more Read Replica
H. Backup: web and app servers backed up weekly via AM Is, database backed up via DB snapshots.
I. Web servers: store read-only data in 53, and copy from 53 to root volume at boot tim
J. App servers: share state using a combination of DynamoDB and IP unicas
K. Database: use RDS with multi-AZ deployment and one or more Read Replica
L. Backup: web and app servers backed up weekly viaAM Is, database backed up via DB snapshots.
M. Web servers: store read-only data in 53, and copy from 53 to root volume at boot tim
N. App servers: share state using a combination of DynamoDB and IP unicas
O. Database: use RDS with multi-AZ deploymen
P. Backup: web and app sewers backed up weekly via ANI Is, database backed up via DB snapshots.

Answer: C

Explanation:
Amazon RDS Multi-AZ deployments provide enhanced availability and durability for Database (DB) Instances, making them a natural fit for production database workloads. When you provision a Multi-AZ DB Instance, Amazon RDS automatically creates a primary DB Instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). Each AZ runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. In case of an infrastructure failure (for example, instance hardware failure, storage failure, or network disruption), Amazon RDS performs an automatic failover to the standby, so that you can resume database operations as soon as the failover is complete. Since the endpoint for your DB Instance remains the same after a failover, your application can resume database operation without the need for manual administrative intervention.
Benefits
Enhanced Durability
Multi-AZ deployments for the MySQL, Oracle, and PostgreSQL engines utilize synchronous physical replication to keep data on the standby up-to-date with the primary. Multi-AZ deployments for the SQL Server engine use synchronous logical replication to achieve the same result, employing SQL
Server-native 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 OS patching or DB Instance scaling, these operations are applied first on the standby, prior to the automatic failover. As a result, your availability impact is, again, only the time required for automatic fail over to complete.
Unlike Single-AZ deployments, 1/0 actMty is not suspended on your primary during backup for MuIti-AZ deployments for the MySOL, 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 Mu|ti-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.
No Administrative Intervention
DB Instance failover is fully automatic and requires no administrative intervention. Amazon RDS monitors the health of your primary and standbys, and initiates a failover automatically in response to a variety of failure conditions.
Failover conditions
Amazon RDS detects and automatically recovers from the most common failure scenarios for Multi-AZ deployments so that you can resume database operations as quickly as possible without administrative intervention. Amazon RDS automatically performs a failover in the event of any of the following:
Loss of availability in primary Availability Zone Loss of network connectMty to primary Compute unit failure on primary
Storage failure on primary
Note: When operations such as DB Instance scaling or system upgrades like OS patching are initiated for Multi-AZ deployments, for enhanced availability, they are applied first on the standby prior to an automatic failover. As a result, your availability impact is limited only to the time required for automatic failover to complete.
Note that Amazon RDS Multi-AZ deployments do not failover automatically in response to database operations such as long running queries, deadlocks or database corruption errors.

NEW QUESTION 153
Company B is launching a new game app for mobile devices. Users will log into the game using their existing social media account to streamline data capture.

Company B would like to directly save player data and scoring information from the mobile app to a DynamoDS table named Score Data
When a user saves their game the progress data will be stored to the Game state 53 bucket. What is the best approach for storing data to DynamoDB and 53?

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

Answer: B

Explanation:
Web Identity Federation
Imagine that you are creating a mobile app that accesses AWS resources, such as a game that runs on a mobile device and stores player and score information using Amazon 53 and DynamoDB. When you write such an app, you'll make requests to AWS services that must be signed with an AWS access key. However, we strongly recommend that you do not embed or distribute long-term AWS credentials with apps that a user downloads to a device, even in an encrypted store. Instead, build your app so that it requests temporary AWS security credentials dynamically when needed using web identity federation. The supplied temporary credentials map to an AWS role that has only the permissions needed to perform
the tasks required by the mobile app.
With web identity federation, you don't need to create custom sign-in code or manage your own user identities. Instead, users of your app can sign in using a well-known identity provider (IdP) - such as Login with Amazon, Facebook, Google, or any other OpenID Connect (OIDC)-compatible IdP, receive an authentication token, and then exchange that token for temporary security credentials in AWS that map to an IAM role with permissions to use the resources in your AWS account. Using an IdP helps you keep your AWS account secure, because you don't have to embed and distribute longterm security credentials with your application.
For most scenarios, we recommend that you use Amazon Cognito because it acts as an identity broker and does much of the federation work for you. For details, see the following section, Using Amazon Cognito for Mobile Apps.
If you don't use Amazon Cognito, then you must write code that interacts with a web IdP (Login with Amazon, Facebook, Google, or any other OIDC-compatible IdP) and then calls the Assume Role With Web Identity API to trade the authentication token you get from those IdPs for AWS temporary security credentials. If you have already used this approach for existing apps, you can continue to use it.
Using Amazon Cognito for Nlobile Apps
The preferred way to use web identity federation is to use Amazon Cognito. For example, Adele the developer is building a game for a mobile device where user data such as scores and profiles is stored in Amazon 53 and Amazon DynamoDB. Adele could also store this data locally on the device and use Amazon Cognito to keep it synchronized across devices. She knows that for security and maintenance reasons, long-term AWS security credentials should not be distributed with the game. She also knows that the game might have a large number of users. For all of these reasons, she does not want to create new user identities in IAM for each player. Instead, she builds the game so that users can sign in using an identity that they've already established with a well-known identity provider, such as Login with Amazon, Facebook, Google, or any OpenID Connect {OIDC)-compatible identity provider.
Her game can take advantage of the authentication mechanism from one of these providers to validate the user's identity.
To enable the mobile app to access her AWS resources, Adele first registers for a developer 10 with her chosen IdPs. She also configures the application with each of these providers. In her AWS account that contains the Amazon 53 bucket and DynamoDB table for the game, Adele uses Amazon Cognito to create IAM roles that precisely define permissions that the game needs. If she is using an OIDC IdP, she also creates an IAM OIDC identity provider entity to establish t rust between her AWS account and the IdP.
In the app's code, Adele calls the sign-in interface for the IdP that she configured previously. The IdP handles all the details of letting the user sign in, and the app gets an OAuth access token or OIDC ID token from the provider. Adele's app can trade this authentication information for a set of temporary security credentials that consist of an AWS access key 10, a secret access key, and a session token.
The app can then use these credentials to access web services offered by AWS. The app is limited to the permissions that are defined in the role that it assumes.
The following figure shows a simplified flow for how this might work, using Login with Amazon as the IdP.
For Step 2, the app can also use Facebook, Google, or any OIDC-compatible identity provider, but that's not shown here.
Sample workflow using Amazon Cognito to federate users for a mobile application

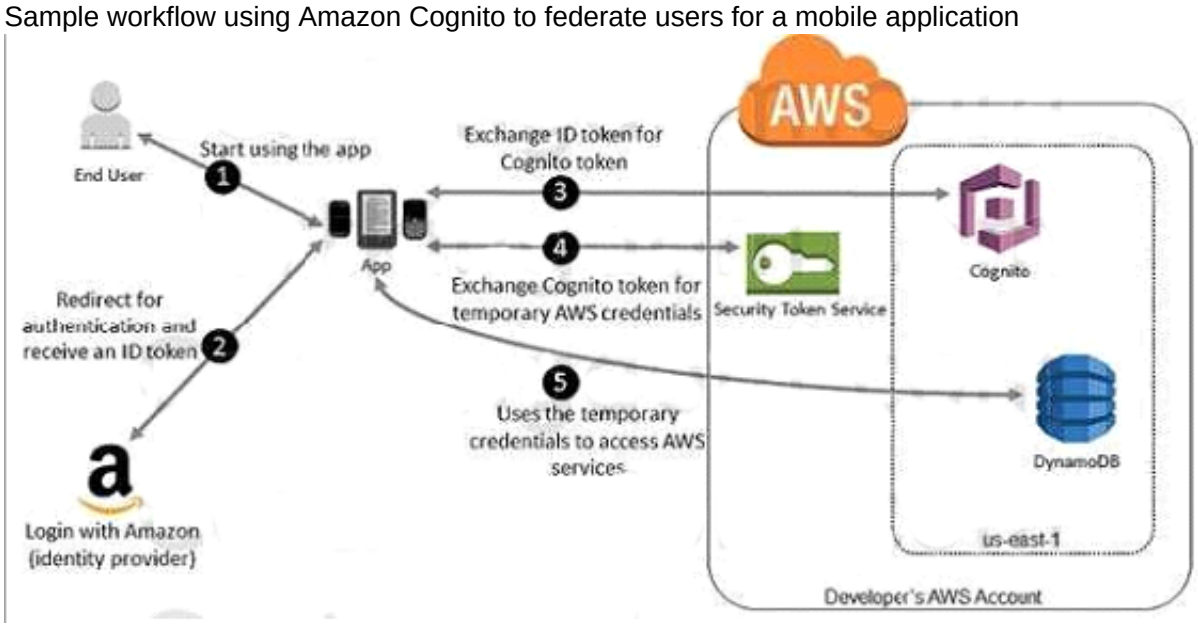

A customer starts your app on a mobile device. The app asks the user to sign in. The app uses Login with Amazon resources to accept the user's credentials.
The app uses Cognito APIs to exchange the Login with Amazon 10 token for a Cognito token. The app requests temporary security credentials from AWS STS, passing the Cognito token.
The temporary security credentials can be used by the app to access any AWS resources required by the app to operate. The role associated with the temporary security credentials and its assigned policies determines what can be accessed.
Use the following process to configure your app to use Amazon Cognito to authenticate users and give your app access to AWS resources. For specific steps to accomplish this scenario, consult the documentation for Amazon Cognito.
(Optional) Sign up as a developer with Login with Amazon, Facebook, Google, or any other OpenID Connect (OIDC}-compatible identity provider and configure one or more apps with the provider. This step is optional because Amazon Cognito also supports unauthenticated (guest) access for your users.
Go to Amazon Cognito in the AWS IV|anagement Console. Use the Amazon Cognito wizard to create an identity pool, which is a container that Amazon Cognito uses to keep end user identities organized for your apps. You can share identity pools between apps. When you set up an identity pool, Amazon Cognito creates one or two IAM roles (one for authenticated identities, and one for unauthenticated "guest" identities) that define permissions for Amazon Cognito users.
Download and integrate the AWS SDK for iOS or the AWS SDK for Android with your app, and import the files required to use Amazon Cognito.
Create an instance of the Amazon Cognito credentials provider, passing the identity pool ID, your AWS account number, and the Amazon Resource Name (ARN) of the ro les that you associated with the identity pool. The Amazon Cognito wizard in the AWS Management Console provides sample code to help you get

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

NEW QUESTION 157
Your company plans to host a large donation website on Amazon Web Services (AWS). You anticipate a large and undetermined amount of traffic that will create many database writes. To be certain that you do not drop any writes to a database hosted on AWS. Which service should you use?

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

Answer: B

Explanation:
Amazon Simple Queue Service (Amazon SQS) offers a reliable, highly scalable hosted queue for storing messages as they travel between computers. By using Amazon SQS, developers can simply move data between distributed application components performing different tasks, without losing messages or requiring each component to be always available. Amazon SQS makes it easy to build a distributed, decoupled application, working in close conjunction with the Amazon Elastic Compute Cloud (Amazon EC2) and the other AWS infrastructure web services.
What can I do with Amazon SQS?
Amazon SQS is a web service that gives you access to a message queue that can be used to store messages while waiting for a computer to process them. This allows you to quickly build message queuing applications that can be run on any computer on the internet. Since Amazon SQS is highly scalable and you only pay for what you use, you can start small and grow your application as you wish, with no compromise on performance or reliability. This lets you focus on building sophisticated message-based applications, without worrying about how the messages are stored and managed.
You can use Amazon SQS with software applications in various ways. For example, you can: Integrate Amazon SQS with other AWS infrastructure web services to make applications more reliable and filexible.
Use Amazon SQS to create a queue of work where each message is a task that needs to be completed by a process. One or many computers can read tasks from the queue and perform them. Build a microservices architecture, using queues to connect your microservices.
Keep notifications of significant events in a business process in an Amazon SQS queue. Each event can have a corresponding message in a queue, and applications that need to be aware of the event can read and process the messages.

NEW QUESTION 158
Your company is in the process of developing a next generation pet collar that collects biometric information to assist families with promoting healthy lifestyles for their pets Each collar will push 30kb of biometric data In JSON format every 2 seconds to a collection platform that will process and analyze the data providing health trending information back to the pet owners and veterinarians via a web portal Management has tasked you to architect the collection platform ensuring the following requirements are met.
Provide the ability for real-time analytics of the inbound biometric data Ensure processing of the biometric data is highly durable. Elastic and parallel The results of the analytic processing should be persisted for data mining
Which architecture outlined below win meet the initial requirements for the collection platform?

A. Utilize 53 to collect the inbound sensor data analyze the data from 53 with a daily scheduled Data Pipeline and save the results to a Redshift Cluster.
B. Utilize Amazon Kinesis to collect the inbound sensor data, analyze the data with Kinesis clients and save the results to a Red shift cluster using EMR.
C. Utilize SQS to collect the inbound sensor data analyze the data from SQS with Amazon Kinesis and save the results to a Mcrosoft SQL Server RDS instance.
D. Utilize EMR to collect the inbound sensor data, analyze the data from EUR with Amazon Kinesis and save me results to Dynamo DB.

Answer: B

NEW QUESTION 162
A web design company currently runs several FTP servers that their 250 customers use to upload and download large graphic files They wish to move this system to AWS to make it more scalable, but they wish to maintain customer privacy and Keep costs to a minimum.
What AWS architecture would you recommend?

A. ASK their customers to use an 53 client instead of an FTP clien
B. Create a single 53 bucket Create an IAM user for each customer Put the IAM Users in a Group that has an IAM policy that permits access to sub-directories within the bucket via use of the 'username' Policy variable.
C. Create a single 53 bucket with Reduced Redundancy Storage turned on and ask their customers to use an 53 client instead of an FTP client Create a bucket for each customer with a Bucket Policy that permits access only to that one customer.
D. Create an auto-scaling group of FTP servers with a scaling policy to automatically scale-in when minimum network traffic on the auto-scaling group is below a given threshol
E. Load a central list of ftp users from 53 as part of the user Data startup script on each Instance.
F. Create a single 53 bucket with Requester Pays turned on and ask their customers to use an 53 client instead of an FTP client Create a bucket tor each customer with a Bucket Policy that permits access only to that one customer.

Answer: A

NEW QUESTION 165
You would like to create a mirror image of your production environment in another region for disaster recovery purposes. Which of the following AWS resources do not need to be recreated in the second region? (Choose 2 answers)

A. Route 53 Record Sets
B. IM Roles
C. Elastic IP Addresses (EIP)
D. EC2 Key Pairs
E. Launch configurations
F. Security Groups

Answer: AC

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

NEW QUESTION 166
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 Multi-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 167
An International company has deployed a multi-tier web application that relies on DynamoDB in a single region For regulatory reasons they need disaster recovery capability In a separate region with a Recovery Time Objective of 2 hours and a Recovery Point Objective of 24 hours They should synchronize their data on a regular basis and be able to provision me web application rapidly using CIoudFormation.
The objective is to minimize changes to the existing web application, control the throughput of DynamoDB used for the synchronization of data and synchronize only the modified elements.
Which design would you choose to meet these requirements?

A. Use AWS data Pipeline to schedule a DynamoDB cross region copy once a da
B. create a Last updated' attribute in your DynamoDB table that would represent the timestamp of the last update and use it as a filter.
C. Use EMR and write a custom script to retrieve data from DynamoDB in the current region using a SCAN operation and push it to Dynamo DB in the second region.
D. Use AWS data Pipeline to schedule an export of the DynamoDB table to 53 in the current region once a day then schedule another task immediately after it that will import data from 53 to DynamoDB in the other region.
E. Send also each Ante into an SOS queue in me second region; use an auto-scaling group behind the SOS queue to replay the write in the second region.

Answer: A

NEW QUESTION 170
An ERP application is deployed across multiple AZs in a single region. In the event of failure, the Recovery Time Objective (RTO) must be less than 3 hours, and the Recovery Point Objective (RPO) must be 15 minutes the customer realizes that data corruption occurred roughly 1.5 hours ago.
What DR strategy could be used to achieve this RTO and RPO in the event of this kind of failure?

A. Take hourly DB backups to 53, with transaction logs stored in 53 every 5 minutes.
B. Use synchronous database master-slave replication between two availability zones.
C. Take hourly DB backups to EC2 Instance store volumes with transaction logs stored In 53 every 5 minutes.
D. Take 15 minute DB backups stored In Glacier with transaction logs stored in 53 every 5 minute

Answer: A

NEW QUESTION 171
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 174
Your company previously configured a heavily used, dynamically routed VPN connection between your on-premises data center and AWS. You recently provisioned a DirectConnect connection and would like to start using the new connection. After configuring DirectConnect settings in the AWS Console, which of the following options win provide the most seamless transition for your users?

A. Delete your existing VPN connection to avoid routing loops configure your DirectConnect router with the appropriate settings and verity network traffic is leveraging DirectConnect.
B. Configure your DirectConnect router with a higher 8GP priority man your VPN router, verify network traffic is leveraging Directconnect and then delete your existing VPN connection.
C. Update your VPC route tables to point to the DirectConnect connection configure your DirectConnect router with the appropriate settings verify network traffic is leveraging DirectConnect and then delete the VPN connection.
D. Configure your DirectConnect router, update your VPC route tables to point to the DirectConnect connection, configure your VPN connection with a higher BGP point
E. And verify network traffic is leveraging the DirectConnect connection.

Answer: D

NEW QUESTION 179
A web company is looking to implement an external payment service into their highly available application deployed in a VPC Their application EC2 instances are behind a public lacing ELB Auto scaling is used to add additional instances as traffic increases under normal load the application runs 2 instances in the Auto Scaling group but at peak it can scale 3x in size. The application instances need to communicate with the payment service over the Internet which requires whitelisting of all public IP addresses used to communicate with it. A maximum of 4 whitelisting IP addresses are allowed at a time and can be added through an API.
How should they architect their solution?

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

Answer: D

NEW QUESTION 183
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 186
You have deployed a three-tier web application in a VPC with a CIOR block of 10 0 0 0/ 28 You initially deploy two web servers, two application sewers, two database sewers and one NAT instance tor a total of seven EC2 instances The web. Application and database servers are deployed across two availability zones

(AZs). You also deploy an ELB in front of the two web sewers, and use Route53 for DN5 Web (raffile gradually increases in the first few days following the deployment, so you attempt to double the number of instances in each tier of the application to handle the new load unfortunately some of these new instances fail to launch.
Which of the following could De the root caused? (Choose 2 answers)

A. AW5 resewes the first and the last private IP address in each subnet's CIDR block so you do not have enough addresses left to launch all of the new EC2 instances
B. The Internet Gateway (IGW) of your VPC has scaled-up, adding more instances to handle the traffic spike, reducing the number of available private IP addresses for new instance launches
C. The ELB has scaled-up, adding more instances to handle the traffic spike, reducing the number of available private IP addresses for new instance launches
D. AW5 reserves one IP address in each subnet's CIDR block for Route53 so you do not have enough addresses left to launch all of the new EC2 instances
E. AW5 reserves the first four and the last IP address in each subnet's CIDR block so you do not haveenough addresses left to launch all of the new EC2 instances

Answer: CE

NEW QUESTION 188
You've been brought in as solutions architect to assist an enterprise customer with their migration of an e-commerce platform to Amazon Virtual Private Cloud (VPC) The previous architect has already deployed a 3-tier VPC, The configuration is as follows:
VPC: vpc-2f8bc447 IGW: igw-2d8bc445 NACL: ad-208bc448
5ubnets and Route Tables: Web sewers: subnet-258bc44d
Application servers: subnet-248bc44c Database sewers: subnet-9189c6f9 Route Tables:
rrb-218bc449 rtb-238bc44b Associations:
subnet-258bc44d : rtb-218bc449 subnet-248bc44c : rtb-238bc44b subnet-9189c6f9 : rtb-238bc44b
You are now ready to begin deploying EC2 instances into the VPC Web servers must have direct access to the internet Application and database sewers cannot have direct access to the internet.
Which configuration below will allow you the ability to remotely administer your application and database servers, as well as allow these sewers to retrieve updates from the Internet?

A. Create a bastion and NAT instance in subnet-258bc44d, and add a route from rtb- 238bc44b to the NAT instance.
B. Add a route from rtb-238bc44b to igw-2d8bc445 and add a bastion and NAT instance within subnet-248bc44c.
C. Create a bastion and NAT instance in subnet-248bc44c, and add a route from rtb- 238bc44b to subneb258bc44d.
D. Create a bastion and NAT instance in subnet-258bc44d, add a route from rtb-238bc44b to Igw- 2d8bc445, and a new NACL that allows access between subnet-258bc44d and subnet -248bc44c.

Answer: A

NEW QUESTION 192
A newspaper organization has a on-premises application which allows the public to search its back catalogue and retrieve indMdual newspaper pages via a website written in Java They have scanned the old newspapers into JPEGs (approx 17TB) and used Optical Character Recognition (OCR) to populate a commercial search product. The hosting platform and software are now end of life and the organization wants to migrate Its archive to AW5 and produce a cost efficient architecture and still be designed for availability and durability. Which is the most appropriate?

A. Use 53 with reduced redundancy Io store and serve the scanned files, install the commercial search application on EC2 Instances and configure with auto-scaling and an Elastic Load Balancer.
B. Model the environment using CIoudFormation use an EC2 instance running Apache webserver and an open source search application, stripe multiple standard EB5 volumes together to store the JPEGs and search index.
C. Use 53 with standard redundancy to store and serve the scanned files, use CIoud5earch for queryprocessing, and use Elastic Beanstalk to host the website across multiple availability zones.
D. Use a single-AZ RD5 My5QL instance Io store the search index 33d the JPEG images use an EC2 instance to serve the website and translate user queries into 5QL.
E. Use a CIoudFront download distribution to serve the JPEGs to the end users and Install the current commercial search product, along with a Java Container Tor the website on EC2 instances and use Route53 with DNS round-robin.

Answer: C

Explanation:
There is no such thing as "NIost appropriate" without knowing all your goals. I find your scenarios very fuzzy, since you can obviously mix-n-match between them. I think you should decide by layers instead: Load Balancer Layer: ELB or just DNS, or roll-your-own. (Using DNS+EIPs is slightly cheaper, but less reliable than ELB.)
Storage Layer for 17TB of Images: This is the perfect use case for 53. Off-load all the web requests directly to the relevant JPEGs in 53. Your EC2 boxes just generate links to them.
If your app already serves it's own images (not links to images), you might start with EFS. But more than likely, you can just setup a web server to re-write or re-direct all JPEG links to 53 pretty easily.
If you use 53, don't serve directly from the bucket- Serve via a CNAME in domain you control. That way, you can switch in C|oudFront easily.
EBS will be way more expensive, and you'II need 2x the drives if you need 2 boxes. Yuck. Consider a smaller storage format. For example, JPEG200 or WebP or other tools might make for smaller images. There is also the DejaVu format from a while back.
Cache Layer: Adding Cloud Front in front of 53 will help people on the other side of the world-- well, possibly. Typical archives follow a power law. The long tail of requests means that most JPEGs won't be requested enough to be in the cache. So you are only speeding up the most popular objects. You can always wait, and switch in CF later after you know your costs better. (In some cases, it can actually lower costs.)
You can also put CIoudFront in front of your app, since your archive search results should be fairly static. This will also allow you to run with a smaller instance type, since CF will handle much of the load if you do it right.
Database Layer: A few options:
Use whatever your current server does for now, and replace with something else down the road. Don't under-estimate this approach, sometimes it's better to start now and optimize later.
Use RDS to run MySQL/ Postgres
I'm not as familiar with EIasticSearch I Cloudsearch, but obviously Cloudsearch will be less maintenance+setup.
App Layer:
When creating the app layer from scratch, consider Cloud Formation and/or OpsWorks. It's extra stuff to learn, but helps down the road.
Java+ Tomcat is right up the alley of E|asticBeanstaIk. (Basically EC2 + Autoscale + ELB).
Preventing Abuse: When you put something in a public 53 bucket, people will hot-link it from their web pages. If you want to prevent that, your app on the EC2 box

can generate signed links to 53 that expire in a few hours. Now everyone will be forced to go thru the app, and the app can apply rate limiting, etc. Saving money:
If you don't mind having downtime:
run everything in one AZ (both DBs and EC2s). You can always add servers and AZs down the road, as long as it's architected to be stateless. In fact, you should use multiple regions if you want it to be really robust.
use Reduced Redundancy in 53 to save a few hundred bucks per month (Someone will have to "go fix it" every time it breaks, including having an off-line copy to repair 53.)
Buy Reserved Instances on your EC2 boxes to make them cheaper. (Start with the RI market and buy a partially used one to get started.) It's just a coupon saying "if you run this type of box in this AZ, you will save on the per-hour costs." You can get 1/2 to 1/3 off easily.
Rewrite the application to use less memory and CPU -that way you can run on fewer/ smaller boxes. (Nlay or may not be worth the investment.)
If your app will be used very infrequently, you will save a lot of money by using Lambda. I'd be worried that it would be quite slow if you tried to run a Java application on it though ..
We're missing some information like load, latency expectations from search, indexing speed, size of the search index, etc. But with what you've given us, I would go with 53 as the storage for the files (53 rocks. It is really, really awesome). If you're stuck with the commercial search application, then on EC2 instances with autoscaling and an ELB. If you are allowed an alternative search engine, Elasticsearch is probably your best bet. I'd run it on EC2 instead ofthe AWS Elasticsearch service, as IMHO it's not ready yet. Don't autoscale Elasticsearch automatically though, it'II cause all sorts of issues. I have zero experience with CloudSearch so I can't comment on that. Regardless of which option, I'd use Cloud Formation for all of it.

NEW QUESTION 194
A corporate web application is deployed within an Amazon Virtual Private Cloud (VPC) and is connected to the corporate data center via an IPsec VPN. The application must authenticate against the on-premises LDAP server. After authentication, each logged-in user can only access an Amazon Simple Storage Space (53) keyspace specific to that user.
Which two approaches can satisfy these objectives? (Choose 2 answers)

A. Develop an identity broker that authenticates against IAM security Token service to assume a Lam role in order to get temporary AWS security credentials The application calls the identity broker to get AWS temporary security credentials with access to the appropriate 53 bucket.
B. The application authenticates against LDAP and retrieves the name of an IAM role associated with the use
C. The application then ca Ils the IAM Security Token Service to assume that IAM role The application can use the temporary credentials to access the appropriate 53 bucket.
D. Develop an identity broker that authenticates against LDAP and then calls IAM Security To ken Service to get IAM federated user credentials The application calls the identity broker to get IAM federated user credentials with access to the appropriate 53 bucket.
E. The application authenticates against LDAP the application then calls the AWS identity and Access Management (IAM) Security service to log in to IAM using the LDAP credentials the application can use the IAM temporary credentials to access the appropriate 53 bucket.
F. The application authenticates against IAM Security Token Service using the LDAP credentials the application uses those temporary AWS security credentials to access the appropriate 53 bucket.

Answer: BC

NEW QUESTION 195
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 198
You've been hired to enhance the overall security posture for a very large e-commerce site They have a well architected multi-tier application running in a VPC that uses ELBs in front of both the web and the app
tier with static assets served directly from 53 They are using a combination of RDS and DynamoOB for their dynamic data and then archMng nightly into 53 for further processing with EMR

They are concerned because they found QUESTION able log entries and suspect someone is attempting to gain unauthorized access.
Which approach provides a cost effective scalable mitigation to this kind of attack?

A. Recommend that they lease space at a DirectConnect partner location and establish a IG DirectConnect connection to their vPC they would then establish Internet connectMty into their space, filter the traffic in hardware Web Application Firewall (WAF). And then pass the traffic through the DirectConnect connection into their application running in their VPC,
B. Add previously identified hostile source IPs as an explicit INBOUND DENY NACL to the web tier sub net.
C. Add a WAF tier by creating a new ELB and an AutoScaling group of EC2 Instances running a host based WAF They would redirect Route 53 to resolve to the new WAF tier ELB The WAF tier would thier pass the traffic to the current web tier The web tier Security Groups would be updated to only allow traffic from the WAF tier Security Group
D. Remove all but TLS 1 2 from the web tier ELB and enable Advanced Protocol Filtering This will enable the ELB itself to perform WAF functionality.

Answer: C

NEW QUESTION 199
You currently operate a web application In the AWS US-East region The application runs on an autoscaled layer of EC2 instances and an RDS Multi-AZ database Your IT security compliance officer has tasked you to develop a reliable and durable logging solution to track changes made to your EC2.1AM And RDS resources. The solution must ensure the integrity and confidentiality of your log data. Which of these solutions would you recommend?

A. Create a new C|oudTraiI trail with one new 53 bucket to store the logs and with the global services option selected Use IAM roles 53 bucket policies and Multi Factor Authentication (MFA) Delete on the 53 bucket that stores your logs.
B. Create a new CIoudTraiI with one new 53 bucket to store the logs Configure SNS to send log file delivery notifications to your management system Use IAM roles and 53 bucket policies on the 53 bucket mat stores your logs.
C. Create a new CIoudTraiI trail with an existing 53 bucket to store the logs and with the global services option selected Use 53 ACLs and Multi Factor Authentication (MFA) Delete on the 53 bucket that stores your logs.
D. Create three new C|oudTrai| trails with three new 53 buckets to store the logs one for the AWS Management console, one for AWS 5DKs and one for command line tools Use IAM roles and 53 bucket policies on the 53 buckets that store your logs.

Answer: A

NEW QUESTION 204
An administrator is using Amazon CIoudFormation to deploy a three tier web application that consists of a web tier and application tier that will utilize Amazon DynamoDB for storage when creating the CIoudFormation template which of the following would allow the application instance access to the DynamoDB tables without exposing API credentials?

A. Create an Identity and Access Management Role that has the required permissions to read and write from the required DynamoDB table and associate the Role to the application instances by referencing an instance profile.
B. Use the Parameter section in the Cloud Formation template to nave the user input Access and Secret Keys from an already created IAM user that has me permissions required to read and write from the required DynamoDB table.
C. Create an Identity and Access Management Role that has the required permissions to read and write from the required DynamoDB table and reference the Role in the instance profile property of the application instance.
D. Create an identity and Access Management user in the CIoudFormation template that has permissions to read and write from the required DynamoDB table, use the GetAtt function to retrieve the Access and secret keys and pass them to the application instance through user-data.

Answer: C

NEW QUESTION 205
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 206
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 208
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 211
You are designing an intrusion detection prevention (IDS/IPS) solution for a customer web application in a single VPC. You are considering the options for implementing IOS IPS protection for traffic coming from the Internet.
Which of the following options would you consider? (Choose 2 answers)

A. Implement IDS/IPS agents on each Instance running In VPC
B. Configure an instance in each subnet to switch its network interface card to promiscuous mode and analyze network traffic.
C. Implement Elastic Load Balancing with SSL listeners In front of the web applications
D. Implement a reverse proxy layer in front of web servers and configure IDS/ IPS agents on each reverse proxy server.

Answer: BD

NEW QUESTION 215
You have an application running on an EC2 Instance which will allow users to download fl ies from a private 53 bucket using a pre-assigned URL. Before generating the URL the application should verify the existence of the fi Ie in 53.
How should the application use AWS credentials to access the 53 bucket securely?

A. Use the AWS account access Keys the application retrieves the credentials from the source code of the application.
B. Create an IAM user for the application with permissions that allow list access to the 53 bucket launch the instance as the IAM user and retrieve the IAM user's credentials from the EC2 instance user data.
C. Create an IAM role for EC2 that allows list access to objects in the 53 bucke
D. Launch the instance with the role, and retrieve the roIe's credentials from the EC2 Instance metadata
E. Create an IAM user for the application with permissions that allow list access to the 53 bucke
F. The application retrieves the IAM user credentials from a temporary directory with permissions that allow read access only to the application user.

Answer: C

NEW QUESTION 218
A benefits enrollment company is hosting a 3-tier web application running in a VPC on AWS which includes a NAT (Network Address Translation) instance in the public Web tier. There is enough provisioned capacity for the expected workload tor the new fiscal year benefit enrollment period plus some extra overhead Enrollment proceeds nicely for two days and then the web tier becomes unresponsive, upon investigation using CIoudWatch and other monitoring tools it is discovered that there is an extremely large and unanticipated amount of inbound traffic coming from a set of 15 specific IP addresses over port 80 from a country where the benefits company has no customers. The web tier instances are so overloaded that benefit enrollment administrators cannot even SSH into them. Which actMty would be useful in defending against this attack?

A. Create a custom route table associated with the web tier and block the attacking IP addresses from the IGW (Internet Gateway)
B. Change the EIP (Elastic IP Address) of the NAT instance in the web tier subnet and update the Main Route Table with the new EIP
C. Create 15 Security Group rules to block the attacking IP addresses over port 80
D. Create an inbound NACL (Network Access control list) associated with the web tier subnet with deny rules to block the attacking IP addresses

Answer: D

Explanation:
Use AWS Identity and Access Management (IAM) to control who in your organization has permission to create and manage security groups and network ACLs (NACL). Isolate the responsibilities and roles for
better defense. For example, you can give only your network administrators or security ad min the permission to manage the security groups and restrict other roles.

NEW QUESTION 223
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 227
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

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

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