AWS-Solution-Architect-Associate Dumps

Amazon AWS Certified Solutions Architect - Associate

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

NEW QUESTION 1
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 2
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 EIastiCache?

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 3
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/general/latest/gr/aws-sec-cred-types.html

NEW QUESTION 4
You need to import several hundred megabytes of data from a local Oracle database to an Amazon RDS DB instance. What does AWS recommend you use to accomplish this?

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

Answer: C

Explanation:
How you import data into an Amazon RDS DB instance depends on the amount of data you have and the number and variety of database objects in your database.
For example, you can use Oracle SQL Developer to import a simple, 20 MB database; you want to use Oracle Data Pump to import complex databases or databases that are several hundred megabytes or several terabytes in size.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Oracle.Procedural.Importing.html

NEW QUESTION 5
A user wants to use an EBS-backed Amazon EC2 instance for a temporary job. Based on the input data, the job is most likely to finish within a week. Which of the following steps should be followed to terminate the instance automatically once the job is finished?

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

Answer: C

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

NEW QUESTION 6
Which of the following AWS CLI commands is syntactically incorrect?
1. $ aws ec2 describe-instances
2. $ aws ec2 start-instances --instance-ids i-1348636c
3. $ aws sns publish --topic-arn arn:aws:sns:us-east-1:546419318123:OperationsError -message "Script Failure"
4. $ aws sqs receive-message --queue-url https://queue.amazonaws.com/546419318123/Test

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

Answer: A

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

NEW QUESTION 7
You have been doing a lot of testing of your VPC Network by deliberately failing EC2 instances to test whether instances are failing over properly. Your customer who will be paying the AWS bill for all this asks you if he being charged for all these instances. You try to explain to him how the billing works on EC2 instances to the best of your knowledge. What would be an appropriate response to give to the customer
in regards to this?

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

Answer: C

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

NEW QUESTION 8
What does the following policy for Amazon EC2 do?
{
"Statement":[{
"Effect":"AI|ow", "Action":"ec2:Describe*", "Resource":"*"
II
}

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

Answer: A

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

NEW QUESTION 9
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 10
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 10
You need to measure the performance of your EBS volumes as they seem to be under performing. You have come up with a measurement of 1,024 KB I/O but your colleague tells you that EBS volume performance is measured in IOPS. How many IOPS is equal to 1,024 KB I/O?

A. 16
B. 256
C. 8
D. 4

Answer: D

Explanation:
Several factors can affect the performance of Amazon EBS volumes, such as instance configuration, I/O characteristics, workload demand, and storage configuration.
IOPS are input/output operations per second. Amazon EBS measures each I/O operation per second
(that is 256 KB or smaller) as one IOPS. I/O operations that are larger than 256 KB are counted in 256 KB capacity units.
For example, a 1,024 KB I/O operation would count as 4 IOPS.
When you provision a 4,000 IOPS volume and attach it to an EBS-optimized instance that can provide the necessary bandwidth, you can transfer up to 4,000 chunks of data per second (provided that the I/O does not exceed the 128 MB/s per volume throughput limit of General Purpose (SSD) and Provisioned IOPS (SSD) volumes).
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSPerformance.htmI

NEW QUESTION 12
A user has created an application which will be hosted on EC2. The application makes calls to DynamoDB to fetch certain data. The application is using the DynamoDB SDK to connect with from the EC2 instance. Which of the below mentioned statements is true with respect to the best practice for security in this scenario?

A. The user should create an IAM user with DynamoDB access and use its credentials within the application to connect with DynamoDB
B. The user should attach an IAM role with DynamoDB access to the EC2 instance
C. The user should create an IAM role, which has EC2 access so that it will allow deploying the application
D. The user should create an IAM user with DynamoDB and EC2 acces
E. Attach the user with the application so that it does not use the root account credentials

Answer: B

Explanation:
With AWS IAM a user is creating an application which runs on an EC2 instance and makes requests to
AWS, such as DynamoDB or S3 calls. Here it is recommended that the user should not create an IAM user and pass the user's credentials to the application or embed those credentials inside the application. Instead, the user should use roles for EC2 and give that role access to DynamoDB /S3. When the roles are attached to EC2, it will give temporary security credentials to the application hosted on that EC2, to connect with DynamoDB / S3.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_WorkingWithGroupsAndUsers.html

NEW QUESTION 16
You are building a system to distribute confidential documents to employees. Using CIoudFront, what method could be used to serve content that is stored in S3, but not publically accessible from S3 directly?

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

Answer: D

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

NEW QUESTION 21
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 23
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 26
A user is running a batch process which runs for 1 hour every day. Which of the below mentioned options is the right instance type and costing model in this case if the user performs the same task for the whole year?

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

Answer: A

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

NEW QUESTION 30
You have just set up a large site for a client which involved a huge database which you set up with Amazon RDS to run as a Multi-AZ deployment. You now start to worry about what will happen if the database instance fails. Which statement best describes how this database will function if there is a database failure?

A. Updates to your DB Instance are synchronously replicated across Availability Zones to the standby in order to keep both in sync and protect your latest database updates against DB Instance failure.
B. Your database will not resume operation without manual administrative intervention.
C. Updates to your DB Instance are asynchronously replicated across Availability Zones to the standby in order to keep both in sync and protect your latest database updates against DB Instance failure.
D. Updates to your DB Instance are synchronously replicated across S3 to the standby in order to keep both in sync and protect your latest database updates against DB Instance failure.

Answer: A

Explanation:
Amazon Relational Database Service (Amazon RDS) is a managed service that makes it easy to set up, operate, and scale a relational database in the cloud. It

provides cost-efficient and resizable capacity, while managing time-consuming database administration tasks, freeing you up to focus on your applications and business.

When you create or modify your DB Instance to run as a Multi-AZ deployment, Amazon RDS automatically provisions and maintains a synchronous "standby" replica in a different Availability Zone. Updates to your DB Instance are synchronously replicated across Availability Zones to the standby in order to keep both in sync and protect your latest database updates against DB Instance failure.

During certain types of planned maintenance, or in the unlikely event of DB Instance failure or Availability Zone failure, Amazon RDS will automatically failover to the standby so that you can resume database writes and reads as soon as the standby is promoted. Since the name record for your DB Instance remains the same, you application can resume database operation without the need for manual administrative intervention. With Multi-AZ deployments, replication is transparent: you do not interact directly with the standby, and it cannot be used to serve read traffic. If you are using Amazon RDS for MySQL and are looking to scale read traffic beyond the capacity constraints of a single DB Instance, you can deploy one or more Read Replicas.

Reference: http://aws.amazon.com/rds/faqs/

NEW QUESTION 31
Name the disk storage supported by Amazon Elastic Compute Cloud (EC2).

A. None of these
B. Amazon AppStream store
C. Amazon SNS store
D. Amazon Instance Store

Answer: D

Explanation:
Amazon EC2 supports the following storage options: Amazon Elastic Block Store (Amazon EBS) Amazon EC2 Instance Store Amazon Simple Storage Service (Amazon S3)
Reference: http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/Storage.html

NEW QUESTION 34
You need to set up security for your VPC and you know that Amazon VPC provides two features that you can use to increase security for your VPC: security groups and network access control lists (ACLs). You have already looked into security groups and you are now trying to understand ACLs. Which statement below is incorrect in relation to ACLs?

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

Answer: B

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

NEW QUESTION 37
While creating an Amazon RDS DB, your first task is to set up a DB that controls which IP address or EC2 instance can access your DB Instance.

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

Answer: D

Explanation:
While creating an Amazon RDS DB, your first task is to set up a DB Security Group that controls what IP addresses or EC2 instances have access to your DB Instance.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithSecurityGroups.html

NEW QUESTION 39
You are setting up some EBS volumes for a customer who has requested a setup which includes a RAID (redundant array of inexpensive disks). AWS has some recommendations for RAID setups. Which RAID setup is not recommended for Amazon EBS?

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

Answer: B

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

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

NEW QUESTION 41
You receive the following request from a client to quickly deploy a static website for them, specifically on AWS. The requirements are low-cost, reliable, online storage, and a reliable and cost-effective way to route customers to the website, as well as a way to deliver content with low latency and high data transfer speeds so that visitors to his website don't experience unnecessary delays. What do you think would be the minimum AWS services that could fulfill the client's request?

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

Answer: C

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

NEW QUESTION 46
You are planning and configuring some EBS volumes for an application. In order to get the most performance out of your EBS volumes, you should attach them to an instance with enough to support your volumes.

A. Redundancy
B. Storage
C. Bandwidth
D. Memory

Answer: C

Explanation:
When you plan and configure EBS volumes for your application, it is important to consider the configuration of the instances that you will attach the volumes to. In order to get the most performance out of your EBS volumes, you should attach them to an instance with enough bandwidth to support your volumes, such as an EBS-optimized instance or an instance with 10 Gigabit network connectMty. This is especially important when you use General Purpose (SSD) or Provisioned IOPS (SSD) volumes, or when you stripe multiple volumes together in a RAID configuration.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-ec2-config.htmI

NEW QUESTION 49
How long does an AWS free usage tier EC2 last for?

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

Answer: B

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

NEW QUESTION 50
Which of the following statements is true of tagging an Amazon EC2 resource?

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

Answer: C

Explanation:
You can assign tags only to resources that already exist. You can't terminate, stop, or delete a resource based solely on its tags; you must specify the resource identifier.
Reference: http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/Using_Tags.html

NEW QUESTION 53
You have three Amazon EC2 instances with Elastic IP addresses in the US East (Virginia) region, and you want to distribute requests across all three IPs evenly for users for whom US East (Virginia) is the appropriate region.
How many EC2 instances would be sufficient to distribute requests in other regions?

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

Answer: D

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

NEW QUESTION 57
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 62
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 67
A gaming company comes to you and asks you to build them infrastructure for their site. They are not sure how big they will be as with all start ups they have limited money and big ideas. What they do tell you is that if the game becomes successful, like one of their previous games, it may rapidly grow to millions of users and generate tens (or even hundreds) of thousands of writes and reads per second. After
considering all of this, you decide that they need a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. Which of the following databases do you think would best fit their needs?

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

Answer: A

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

NEW QUESTION 69
You're trying to delete an SSL certificate from the IAM certificate store, and you're getting the message "Certificate: <certificate-id> is being used by CloudFront."
Which of the following statements is probably the reason why you are getting this error?

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

Answer: A

Explanation:
CloudFront is a web service that speeds up distribution of your static and dynamic web content, for example, .html, .css, .php, and image files, to end users.
Every CloudFront web distribution must be associated either with the default CloudFront certificate or with a custom SSL certificate. Before you can delete an SSL certificate, you need to either rotate SSL certificates (replace the current custom SSL certificate with another custom SSL certificate) or revert from using a custom SSL certificate to using the default CloudFront certificate.
Reference: http://docs.aws.amazon.com/AmazonCloudFront/latest/Deve|operGuide/Troubleshooting.htm|

NEW QUESTION 70
How many types of block devices does Amazon EC2 support?

A. 4
B. 5
C. 2
D. 1

Answer: C

Explanation:
Amazon EC2 supports 2 types of block devices. Reference:
http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/block-device-mapping-concepts.html

NEW QUESTION 71
You are setting up some IAM user policies and have also become aware that some services support resource-based permissions, which let you attach policies to the service's resources instead of to IAM users or groups. Which of the below statements is true in regards to resource-level permissions?

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

Answer: D

Explanation:
AWS Identity and Access Management is a web service that enables Amazon Web Services (AWS) customers to manage users and user permissions in AWS. The service is targeted at organizations with multiple users or systems that use AWS products such as Amazon EC2, Amazon RDS, and the AWS Management Console. With IAM, you can centrally manage users, security credentials such as access keys, and permissions that control which AWS resources users can access.
In addition to supporting IAM user policies, some services support resource-based permissions, which let you attach policies to the service's resources instead of to IAM users or groups. Resource-based permissions are supported by Amazon S3, Amazon SNS, and Amazon SQS.
The resource-level permissions service supports IAM policies in which you can specify indMdual resources using Amazon Resource Names (ARNs) in the policy's Resource element.
Some services support resource-level permissions only for some actions.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_SpecificProducts.html

NEW QUESTION 73
Amazon Elastic Load Balancing is used to manage traffic on a fileet of Amazon EC2 instances, distributing traffic to instances across all availability zones within a region. Elastic Load Balancing has all the advantages of an on-premises load balancer, plus several security benefits.
Which of the following is not an advantage of ELB over an on-premise load balancer?

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

Answer: A

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

NEW QUESTION 78
In Amazon CloudFront, if you use Amazon EC2 instances and other custom origins with CloudFront, it is recommended to .

A. not use Elastic Load Balancing

B. restrict Internet communication to private instances while allowing outgoing traffic
C. enable access key rotation for 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 79
What happens to data on an ephemeral volume of an EBS-backed EC2 instance if it is terminated or if it fails?

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

Answer: D

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

NEW QUESTION 81
Content and IV|edia Server is the latest requirement that you need to meet for a client.
The client has been very specific about his requirements such as low latency, high availability, durability, and access control. Potentially there will be millions of views on this server and because of "spiky" usage patterns, operations teams will need to provision static hardware, network, and management resources to support the maximum expected need. The Customer base will be initially low but is expected to grow and become more geographically distributed.
Which of the following would be a good solution for content distribution?

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

Answer: D

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

NEW QUESTION 82
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 85
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 88
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 92
Which DNS name can only be resolved within Amazon EC2?

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

Answer: B

Explanation:
Only Internal DNS name can be resolved within Amazon EC2. Reference:
http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/using-instance-addressing.html

NEW QUESTION 94
You want to establish a dedicated network connection from your premises to AWS in order to save money by transferring data directly to AWS rather than through your internet service provider. You are sure there must be some other benefits beyond cost savings. Which of the following would not be considered a benefit if you were to establish such a connection?

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

Answer: D

Explanation:
AWS Direct Connect makes it easy to establish a dedicated network connection from your premises to AWS.
Using AWS Direct Connect, you can establish private connectMty between AWS and your datacenter, office, or colocation environment, which in many cases can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than internet-based connections.
You could expect the following benefits if you use AWS Direct Connect. Reduced bandwidth costs
Consistent network performance Compatibility with all AWS services Private connectMty to your Amazon VPC Elasticity
Simplicity
Reference: http://aws.amazon.com/directconnect/

NEW QUESTION 97
A user has launched an EC2 instance. The instance got terminated as soon as it was launched. Which of the below mentioned options is not a possible reason for this?

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

Answer: D

Explanation:

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

NEW QUESTION 101
A user has created an ELB with Auto Scaling. Which of the below mentioned offerings from ELB helps the
user to stop sending new requests traffic from the load balancer to the EC2 instance when the instance is being deregistered while continuing in-flight requests?

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

Answer: D

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

NEW QUESTION 105
You have some very sensitive data stored on AWS S3 and want to try every possible alternative to keeping it secure in regards to access control. What are the mechanisms available for access control on AWS S3?

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

Answer: A

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

NEW QUESTION 106
Which one of the following can't be used as an origin server with Amazon CIoudFront?

A. A web server running in your infrastructure
B. Amazon S3
C. Amazon Glacier
D. A web server running on Amazon EC2 instances

Answer: C

Explanation:
Amazon CIoudFront is designed to work with Amazon S3 as your origin server, customers can also use Amazon C|oudFront with origin sewers running on Amazon EC2 instances or with any other custom origin.
Reference: http://docs.aws.amazon.com/AmazonCIoudFront/latest/DeveIoperGuide/distribution-web.html

NEW QUESTION 110
You have been asked to set up monitoring of your network and you have decided that Cloudwatch would be the best service to use. Amazon CIoudWatch monitors your Amazon Web Services (AWS) resources and the applications you run on AWS in real-time. You can use CIoudWatch to collect and track metrics, which are the variables you want to measure for your resources and applications. Which of the following items listed can AWS Cloudwatch monitor?

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

Answer: B

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

NEW QUESTION 113
How can you apply more than 100 rules to an Amazon EC2-Classic?

A. By adding more security groups
B. You need to create a default security group specifying your required rules if you need to use more than 100 rules per security group.
C. By default the Amazon EC2 security groups support 500 rules.
D. You can't add more than 100 rules to security groups for an Amazon EC2 instanc

Answer: D

Explanation:
In EC2-Classic, you can associate an instance with up to 500 security groups and add up to 100 rules to a security group.
Reference: http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/using-network-security.htmI

NEW QUESTION 118
Which of the following statements is NOT true about using Elastic IP Address (EIP) in EC2-Classic and EC2-VPC platforms?

A. In the EC2-VPC platform, the Elastic IP Address (EIP) does not remain associated with the instance when you stop it.
B. In the EC2-Classic platform, stopping the instance disassociates the Elastic IP Address (EIP) from it.
C. In the EC2-VPC platform, if you have attached a second network interface to an instance, when you disassociate the Elastic IP Address (EIP) from that instance, a new public IP address is not assigned to the instance automatically; you'II have to associate an EIP with it manually.
D. In the EC2-Classic platform, if you disassociate an Elastic IP Address (EIP) from the instance, the instance is automatically assigned a new public IP address within a few minutes.

Answer: A

Explanation:
In the EC2-Classic platform, when you associate an Elastic IP Address (EIP) with an instance, the instance's current public IP address is released to the EC2-Classic public IP address pool. If you disassociate an EIP from the instance, the instance is automatically assigned a new public IP address within a few minutes. In addition, stopping the instance also disassociates the EIP from it.
But in the EC2-VPC platform, when you associate an EIP with an instance in a default Virtual Private Cloud (VPC), or an instance in which you assigned a public IP to the eth0 network interface during launch, its current public IP address is released to the EC2-VPC public IP address pool. If you disassociate an EIP from the instance, the instance is automatically assigned a new public IP address within a few minutes. However, if you have attached a second network interface to the instance, the instance is not automatically assigned a new public IP address; you'II have to associate an EIP with it manually. The EIP remains associated with the instance when you stop it.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-ip-addresses-eip.htmI

NEW QUESTION 123
An application hosted at the EC2 instance receives an HTTP request from ELB. The same request has an X-Forvvarded-For header, which has three IP addresses. Which system's IP will be a part of this header?

A. Previous Request IP address.
B. Client IP address.
C. All of the answers listed here.
D. Load Balancer IP addres

Answer: C

Explanation:
When a user sends a request to ELB over HTTP/HTTPS, the request header log at the instance will only receive the IP of ELB. This is because ELB is the interceptor between the EC2 instance and the client request. To get the client IP, use the header X-Forvvarded-For in header. The client IP address in the X-Fonzvarded-For request header is followed by the IP addresses of each successive proxy that passes along the request. The last IP address is the IP address that connects to the back-end application instance. e.g. if the HTTP request already has a header when it reaches the Load Balancer, the IP address from which the request came is appended at the end of the header followed by the IP address of the Load Balancer. In such cases, the X-Forvvarded-For request header takes the following form:
X-Fonzvarded-For: clientIPAddress, previousRequestIPAddress, LoadBaIancerIPAddress. Reference:
http://docs.aws.amazon.com/E|asticLoadBaIancing/Iatest/DeveIoperGuide/TerminologyandKeyConcepts. html

NEW QUESTION 127
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 Multi-AZ deployments also extend to planned maintenance and backups.
In the case of system upgrades like OS patching or DB Instance scaling, these operations are applied first on the standby, prior to the automatic failover. As a result, your availability impact is, again, only the time required for automatic fail over to complete.
Unlike Single-AZ deployments, 1/0 actMty is not suspended on your primary during backup for Multi-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 Multi-AZ deployments.
On instance failure in Amazon Aurora deployments, Amazon RDS uses RDS Multi-AZ technology to automate failover to one of up to 15 Amazon Aurora Replicas you have created in any of three Availability Zones. If no Amazon Aurora Replicas have been provisioned, in the case of a failure, Amazon RDS will attempt to create a new Amazon Aurora DB instance for you automatically.
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 128
Your customer wishes to deploy an enterprise application to AWS which will consist of several web servers, several application servers and a small (50GB) Oracle database information is stored, both in the database and the file systems of the various servers. The backup system must support database recovery whole server and whole disk restores, and indMdual file restores with a recovery time of no more than two hours. They have chosen to use RDS Oracle as the database
Which backup architecture will meet these requirements?

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

Answer: A

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

NEW QUESTION 130
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 ElastiCache to store the writes until the writes are committed to the database.
D. Amazon DynamoDB with provisioned write throughput up to the anticipated peak write 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 132
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 137
You are migrating a legacy client-server application to AWS. The application responds to a specific DNS domain (e.g. www.examp|e.com) and has a 2-tier architecture, with multiple application sewers and a database server. Remote clients use TCP to connect to the application sewers. The application servers need to know the IP address of the clients in order to function properly and are currently taking that information from the TCP socket. A Multi-AZ RDS MySQL instance will be used for the database. During the migration you can change the application code, but you have to file a change request.
How would you implement the architecture on AWS in order to maximize scalability and high availability?

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

Answer: D

NEW QUESTION 138
You are the new IT architect in a company that operates a mobile sleep tracking application
When activated at night, the mobile app is sending collected data points of 1 kilobyte every 5 minutes to
your backend
The backend takes care of authenticating the user and writing the data points into an Amazon DynamoDB table.
Every morning, you scan the table to extract and aggregate last night's data on a per user basis, and store the results in Amazon 53.
Users are notified via Amazon 5NI5 mobile push notifications that new data is available, which is parsed and visualized by (The mobile app Currently you have around IOOk users who are mostly based out of North America.
You have been tasked to optimize the architecture of the backend system to lower cost what would you recommend? (Choose 2 answers}

A. Create a new Amazon DynamoDB (able each day and drop the one for the previous day after its data is on Amazon 53.
B. Have the mobile app access Amazon DynamoDB directly instead of J50N files stored on Amazon 53.
C. Introduce an Amazon SQS queue to buffer writes to the Amazon DynamoDB table and reduce provisioned write throughput.
D. Introduce Amazon Elasticache Io cache reads from the Amazon DynamoDB table and reduce provisioned read throughput.
E. Write data directly into an Amazon Redshift cluster replacing both Amazon DynamoDB and Amazon 53.

Answer: BD

NEW QUESTION 140
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 145
You require the ability to analyze a customer's clickstream data on a website so they can do behavioral analysis. Your customer needs to know what sequence of pages and ads their customer clicked on. This data will be used in real time to modify the page layouts as customers click through the site to increase stickiness and advertising click-through. Which option meets the requirements for captioning and analyzing this data?

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

Answer: B

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

NEW QUESTION 146
Your customer is willing to consolidate their log streams (access logs application logs security logs etc.) in one single system. Once consolidated, the customer wants to analyze these logs in real time based on heuristics. From time to time, the customer needs to validate heuristics, which requires going back to data samples extracted from the last 12 hours?
What is the best approach to meet your customer's requirements?

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

Answer: B

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

NEW QUESTION 151
You deployed your company website using Elastic Beanstalk and you enabled log file rotation to 53. An Elastic Map Reduce job is periodically analyzing the logs on 53 to build a usage dashboard that you share with your CIO.
You recently improved overall performance of the website using Cloud Front for dynamic content delivery and your website as the origin.
After this architectural change, the usage dashboard shows that the traffic on your website dropped by an order of magnitude. How do you fix your usage dashboard'?

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

Answer: D

NEW QUESTION 153
A web company is looking to implement an intrusion detection and prevention system into their deployed VPC. This platform should have the ability to scale to thousands of instances running inside of the VPC, How should they architect t heir solution to achieve these goals?

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

Answer: C

NEW QUESTION 154
Your firm has uploaded a large amount of aerial image data to 53 In the past, in your on-premises environment, you used a dedicated group of servers to oaten process this data and used Rabbit MOAn open source messaging system to get job information to the servers. Once processed the data would go to tape and be shipped offsite. Your manager told you to stay with the current design, and leverage AWS archival storage and messaging services to minimize cost. Which is correct?

A. Use SOS for passing job messages use Cloud Watch alarms to terminate EC2 worker instances when they become idl
B. Once data is processed, change the storage class of the 53 objects to Reduced Redundancy Storage.
C. Setup Auto-Scaled workers triggered by queue depth that use spot instances to process messages in SOS Once data is processed,
D. Change the storage class of the 53 objects to Reduced Redundancy Storag
E. Setup Auto-Scaled workers triggered by queue depth that use spot instances to process messages in SOS Once data is processed, change the storage class of the 53 objects to Glacier.

F. Use SNS to pass job messages use Cloud Watch alarms to terminate spot worker instances when they become idl
G. Once data is processed, change the storage class of the 53 object to Glacier.

Answer: D

NEW QUESTION 159
What does RRS stand for when talking about 53?

A. Redundancy Removal System
B. Relational Rights Storage
C. Regional Rights Standard
D. Reduced Redundancy Storage

Answer: D

NEW QUESTION 160
All Amazon EC2 instances are assigned two IP addresses at launch, out of which one can only be reached from within the Amazon EC2 network?

A. Multiple IP address
B. Public IP address
C. Private IP address
D. Elastic I P Address

Answer: C

NEW QUESTION 165
What is the Reduced Redundancy option in Amazon 53?

A. Less redundancy for a lower cost.
B. It doesn't exist in Amazon 53, but in Amazon EBS.
C. It allows you to destroy any copy of your files outside a specific jurisdiction.
D. It doesn't exist at all

Answer: A

NEW QUESTION 166
While creating an Amazon RDS DB, your first task is to set up a DB _ that controls what IP addresses or EC2 instances have access to your DB Instance.

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

Answer: D

NEW QUESTION 169
Every user you create in the IAM system starts with _ _

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

Answer: C

NEW QUESTION 173
What does Amazon EC2 provide?

A. Virtual sewers in the Cloud.
B. A platform to run code (Java, PHP, Python), paying on an hourly basis.
C. Computer Clusters in the Cloud.
D. Physical sewers, remotely managed by the custome

Answer: A

NEW QUESTION 176
Can I control if and when MySQL based RDS Instance is upgraded to new supported versions?

A. No
B. Only in VPC
C. Yes

Answer: C

NEW QUESTION 181

What does Amazon Elastic Beanstalk provide?

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

Answer: B

NEW QUESTION 184
What will be the status of the snapshot until the snapshot is complete.

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

Answer: D

NEW QUESTION 187
While creating the snapshots using the command line tools, which command should I be using?

A. ec2-deploy-snapshot
B. ec2-fresh-snapshot
C. ec2-create-snapshot
D. ec2-new-snapshot

Answer: C

NEW QUESTION 190
Groups can't _.

A. be nested more than 3 levels
B. be nested at all
C. be nested more than 4 levels
D. be nested more than 2 levels

Answer: B

NEW QUESTION 191
What are the two types of licensing options available for using Amazon RDS for Oracle?

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

Answer: B

NEW QUESTION 196
Is creating a Read Replica of another Read Replica supported?

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

Answer: D

NEW QUESTION 198
Can a 'user' be associated with multiple AWS accounts?

A. No
B. Yes

Answer: A

NEW QUESTION 202
True or False: Manually created DB Snapshots are deleted after the DB Instance is deleted.

A. TRUE
B. FALSE

Answer: A

NEW QUESTION 203
Can I move a Reserved Instance from one Region to another?

A. No
B. Only if they are moving into GovC|oud
C. Yes
D. Only if they are moving to US East from another region

Answer: A

NEW QUESTION 207
What is Amazon Glacier?

A. You mean Amazon "Iceberg": it's a low-cost storage service.
B. A security tool that allows to "freeze" an EBS volume and perform computer forensics on it.
C. A low-cost storage service that provides secure and durable storage for data archMng and backup.
D. It's a security tool that allows to "freeze" an EC2 instance and perform computer forensics on i

Answer: C

NEW QUESTION 208
Is Federated Storage Engine currently supported by Amazon RDS for MySQL?

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

Answer: B

NEW QUESTION 211
Is there a limit to how many groups a user can be in?

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

Answer: A

NEW QUESTION 213
In the Launch Db Instance Wizard, where can I select the backup and maintenance options?

A. Under DB INSTANCE DETAILS
B. Under REVI EW
C. Under MANAGEMENT OPTIONS
D. Under ENGINE SELECTION

Answer: C

NEW QUESTION 218
IAM provides several policy templates you can use to automatically assign permissions to the groups you create. The _ policy template gives the Admins group permission to access all account resources,
except your AWS account information

A. Read Only Access
B. Power User Access
C. AWS Cloud Formation Read Only Access
D. Administrator Access

Answer: D

NEW QUESTION 221
SQL Sewer _ store log ins and passwords in the master database.

A. can be configured to but by default does not
B. doesn't
C. does

Answer: C

NEW QUESTION 225
What is the minimum charge for the data transferred between Amazon RDS and Amazon EC2 Instances in the same Availability Zone?

A. USD 0.10 per GB

B. No charg
C. It is free.
D. USD 0.02 per GB
E. USD 0.01 per GB

Answer: B

NEW QUESTION 228
Which service enables AWS customers to manage users and permissions in AWS?

A. AWS Access Control Service (ACS}
B. AWS Identity and Access Management (IAM}
C. AWS Identity Manager (AIM}

Answer: B

NEW QUESTION 230
Which Amazon service can I use to define a virtual network that closely resembles a traditional data center?

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

Answer: A

NEW QUESTION 232
MySQL installations default to port _.

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

Answer: A

NEW QUESTION 234
If you have chosen Multi-AZ deployment, in the event of a planned or unplanned outage of your primary DB Instance, Amazon RDS automatically switches to the standby replica. The automatic failover mechanism simply changes the record of the main DB Instance to point to the standby DB Instance.

A. DNAME
B. CNAME
C. TXT
D. MX

Answer: B

NEW QUESTION 239
If I have multiple Read Replicas for my master DB Instance and I promote one of them, what happens to the rest of the Read Replicas?

A. The remaining Read Replicas will still replicate from the older master DB Instance
B. The remaining Read Replicas will be deleted
C. The remaining Read Replicas will be combined to one read replica

Answer: A

NEW QUESTION 243
Select the incorrect statement

A. In Amazon EC2, the private IP addresses only returned to Amazon EC2 when the instance is stopped or terminated
B. In Amazon VPC, an instance retains its private IP addresses when the instance is stopped.
C. In Amazon VPC, an instance does NOT retain its private IP addresses when the instance is stopped.
D. In Amazon EC2, the private IP address is associated exclusive ly with the instance for its lifetime

Answer: C

NEW QUESTION 244
Can I test my DB Instance against a new version before upgrading?

A. Only in VPC
B. No
C. Yes

Answer: C

NEW QUESTION 249
The _ service is targeted at organizations with multiple users or systems that use AWS products such as Amazon EC2, Amazon SimpleDB, and the AWS Management Console.

A. Amazon RDS
B. AWS Integrity Management
C. AWS Identity and Access Management
D. Amazon EMR

Answer: C

NEW QUESTION 250
What does Amazon Route53 provide?

A. A global Content Delivery Network.
B. None of these.
C. A scalable Domain Name System.
D. An SSH endpoint for Amazon EC2.

Answer: C

NEW QUESTION 251
If an Amazon EBS volume is the root device of an instance, can I detach it without stopping the instance?

A. Yes but only if Windows instance
B. No
C. Yes
D. Yes but only if a Linux instance

Answer: B

NEW QUESTION 254
Every user you create in the IAM system starts with _ _

A. full permissions
B. no permissions
C. partial permissions

Answer: B

NEW QUESTION 258
Can I initiate a "forced failover" for my MySQL Multi-AZ DB Instance deployment?

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

Answer: A

NEW QUESTION 260
A group can contain many users. Can a user belong to multiple groups?

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

Answer: A

NEW QUESTION 262
Please select the most correct answer regarding the persistence of the Amazon Instance Store

A. The data on an instance store volume persists only during the life of the associated Amazon EC2 instance
B. The data on an instance store volume is lost when the security group rule of the associated instance is changed.
C. The data on an instance store volume persists even after associated Amazon EC2 instance is deleted

Answer: B

NEW QUESTION 263
Is the SQL Server Audit feature supported in the Amazon RDS SQL Server engine?

A. No
B. Yes

Answer: A

NEW QUESTION 265
My Read Replica appears "stuck" after a Multi-AZ failover and is unable to obtain or apply updates from the source DB Instance. What do I do?

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

Answer: A

NEW QUESTION 266
Select the correct set of steps for exposing the snapshot only to specific AWS accounts

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

Answer: C

NEW QUESTION 270
What does Amazon RDS stand for?

A. Regional Data Sewer.
B. Relational Database Service.
C. Nothing.
D. Regional Database Service.

Answer: B

NEW QUESTION 272
What is the maximum response time for a Business level Premium Support case?

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

Answer: D

NEW QUESTION 273
Is there a method in the IAM system to al low or deny access to a specific instance?

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

Answer: C

NEW QUESTION 278
Does AWS Direct Connect allow you access to all Availabilities Zones within a Region?

A. Depends on the type of connection
B. No
C. Yes
D. Only when there's just one availability zone in a regio
E. If there are more than one, only one availability zone can be accessed directly.

Answer: A

NEW QUESTION 283
True or False: When you add a rule to a DB security group, you do not need to specify port number or protocol.

A. Depends on the ROMS used
B. TRUE
C. FALSE

Answer: B

NEW QUESTION 287
Without -' you must either create multiple AWS accounts-each with its own billing and subscriptions to AWS products-or your employees must share the security credentials of a single AWS account.

A. Amazon RDS
B. Amazon Glacier
C. Amazon EMR
D. Amazon IAM

Answer: D

NEW QUESTION 290
Amazon RDS supports SOAP only through _ _

A. HTTP or HTTPS
B. TCP/IP
C. HTIP
D. HTIPS

Answer: D

NEW QUESTION 291
The Amazon EC2 web service can be accessed using the _ web services messaging protocol. This interface is described by a Web Services Description Language (WSDL) document.

A. SOAP
B. DCOM
C. CORBA
D. XML-RPC

Answer: A

NEW QUESTION 292
When you resize the Amazon RDS DB instance, Amazon RDS will perform the upgrade during the next maintenance window. If you want the upgrade to be performed now, rather than waiting for the maintenance window, specify the _ option.

A. Apply Now
B. Apply Soon
C. Apply This
D. Apply Immediately

Answer: D

NEW QUESTION 295
A _ is a document that provides a formal statement of one or more permissions.

A. policy
B. permission
C. Role
D. resource

Answer: A

NEW QUESTION 300
If I scale the storage capacity provisioned to my DB Instance by mid of a billing month, how will I be charged?

A. You will be charged for the highest storage capacity you have used
B. On a proration basis
C. You will be charged for the lowest storage capacity you have used

Answer: B

NEW QUESTION 302
How can an EBS volume that is currently attached to an EC2 instance be migrated from one Availability Zone to another?

A. Detach the volume and attach it to another EC2 instance in the other AZ.
B. Simply create a new volume in the other AZ and specify the original volume as the source.
C. Create a snapshot of the volume, and create a new volume from the snapshot in the other AZ.
D. Detach the volume, then use the ec2-migrate-voiume command to move it to another A

Answer: C

NEW QUESTION 305
You are using an m1.smaII EC2 Instance with one 300GB EBS volume to host a relational database. You determined that write throughput to the database needs to be increased. Which of the following approaches can help achieve this? Choose 2 answers

A. Use an array of EBS volumes.
B. Enable Multi-AZ mode.
C. Place the instance in an Auto Scaling Groups

D. Add an EBS volume and place into RAID 5.
E. Increase the size of the EC2 Instance.
F. Put the database behind an Elastic Load Balance

Answer: DE

NEW QUESTION 310
You nave multiple Amazon EC2 instances running in a cluster across multiple Availability Zones within the same region. What combination of the following should be used to ensure the highest network performance (packets per second), lowest latency, and lowestjitter? Choose 3 answers

A. Amazon EC2 placement groups
B. Enhanced networking
C. Amazon PV AMI
D. Amazon HVM AM
E. Amazon Linux
F. Amazon VPC

Answer: ABE

NEW QUESTION 313
You have an EC2 Security Group with several running EC2 instances. You change the Security Group rules to allow inbound traffic on a new port and protocol, and launch several new instances in the same Security Group. The new rules apply:

A. Immediately to all instances in the security group.
B. Immediately to the new instances only.
C. Immediately to the new instances, but old instances must be stopped and restarted before the new rules apply.
D. To all instances, but it may take several minutes for old instances to see the change

Answer: A

NEW QUESTION 316
Which technique can be used to integrate AWS IAM (Identity and Access Management) with an on-premise LDAP (Lightweight Directory Access Protocol) directory service?

A. Use an IAM policy that references the LDAP account identifiers and the AWS credentials.
B. Use SANIL (Security Assertion Markup Language) to enable single sign-on between AWS and LDAP.
C. Use AWS Security Token Service from an identity broker to issue short-lived AWS credentials.
D. Use IAM roles to automatically rotate the IAM credentials when LDAP credentials are updated.
E. Use the LDAP credentials to restrict a group of users from launching specific EC2 instance type

Answer: B

NEW QUESTION 317
You launch an Amazon EC2 instance without an assigned AVVS identity and Access Management (IAM) role. Later, you decide that the instance should be running with an IAM role. Which action must you take in order to have a running Amazon EC2 instance with an IAM role assigned to it?

A. Create an image of the instance, and register the image with an IAM role assigned and an Amazon EBS volume mapping.
B. Create a new IAM role with the same permissions as an existing IAM role, and assign it to the running instance.
C. Create an image of the instance, add a new IAM role with the same permissions as the desired IAM role, and deregister the image with the new role assigned.
D. Create an image of the instance, and use this image to launch a new instance with the desired Lam role assigned.

Answer: D

Explanation:
Reference:http://docs.aws.amazon.com/IAM/latest/UserGuide/roles-usingroIe-ec2instanee.htm|

NEW QUESTION 319
What are characteristics of Amazon 53? Choose 2 answers

A. 53 allows you to store objects of virtually unlimited size.
B. 53 offers Provisioned IOPS.
C. 53 allows you to store unlimited amounts of data.
D. 53 should be used to host a relational database.
E. Objects are directly accessible via a UR

Answer: CE

Explanation:
Reference:
http://docs.aws.amazon.com/AmazonCioudFront/latest/DeveIoperGuide/private-contentrestricting-access
-to-s3.htmI

NEW QUESTION 321
A customer needs to capture all client connection information from their load balancer every five minutes. The company wants to use this data for analyzing traffic patterns and troubleshooting their applications. Which of the following options meets the customer requirements?

A. Enable AWS CloudTrail for the load balancer.
B. Enable access logs on the load balancer.
C. Install the Amazon CloudWatch Logs agent on the load balancer.
D. Enable Amazon CloudWatch metrics on the load balance

Answer: A

NEW QUESTION 326
A company needs to deploy virtual desktops to its customers in a virtual private cloud, leveraging existing security controls. Which set of AWS services and features will meet the company's requirements?

A. Virtual Private Network connectio
B. AWS Directory Services, and Classic link
C. Virtual Private Network connectio
D. AWS Di rectory Services, and Amazon Workspaces
E. AWS Directory Service, Amazon Workspaces, and AWS Identity and Access Management
F. Amazon Elastic Compute Cloud, and AWS Identity and Access Management

Answer: C

NEW QUESTION 329
After creating a new IAM user which of the following must be done before they can successfully make API calls?

A. Add a password to the user.
B. Enable Multi-Factor Authentication for the user.
C. Assign a Password Policy to the user.
D. Create a set of Access Keys for the use

Answer: D

Explanation:
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_SettingUpUser.html

NEW QUESTION 333
Which of the following are characteristics of a reserved instance? Choose 3 answers

A. It can be migrated across Availability Zones
B. It is specific to an Amazon Machine Image (AMI)
C. It can be applied to instances launched by Auto Scaling
D. It is specific to an instance Type
E. It can be used to lower Total Cost of Ownership (TCO) of a system

Answer: CDE

NEW QUESTION 334
You are tasked with setting up a Linux bastion host for access to Amazon EC2 instances running in your VPC. Only clients connecting from the corporate external public IP address 72.34.51.100 should have SSH access to the host. Which option will meet the customer requirement?

A. Security Group Inbound Rule: Protocol - TC
B. Port Range- 22, Source 72.34.51. 100/32
C. Security Group Inbound Rule: Protocol - UDP, Port Range- 22, Source 72.34.51.100/32
D. Network ACL Inbound Rule: Protocol - UDP, Port Range- 22, Source 72.34.51.100/32
E. Network ACL Inbound Rule: Protocol - TCP, Port Range-22, Source 72.34.51.100/0

Answer: A

NEW QUESTION 336
You are working with a customer who is using Chef configuration management in their data center. Which service is designed to let the customer leverage existing Chef recipes in AWS?

A. Amazon Simple Workflow Service
B. AWS Elastic Beanstalk
C. AWS CloudFormation
D. AWS OpsWorks

Answer: D

Explanation:
Reference: http://aws.amazon.com/opsworks/

NEW QUESTION 341
In order to optimize performance for a compute cluster that requires low inter-node latency, which of the following feature should you use?

A. Multiple Availability Zones
B. AWS Direct Connect
C. EC2 Dedicated Instances

D. Placement Groups E.VPC private subnets

Answer: D

Explanation:
Reference: http://aws.amazon.com/ec2/faqs/ (enhanced networking)

NEW QUESTION 345
You have an environment that consists of a public subnet using Amazon VPC and 3 instances that are running in this subnet. These three instances can successfully communicate with other hosts on the Internet. You launch a fourth instance in the same subnet, using the same AM and security group configuration you used for the others, but find that this instance cannot be accessed from the internet. What should you do to enable Internet access?

A. Deploy a NAT instance into the public subnet.
B. Assign an Elastic IP address to the fourth instance.
C. Configure a publically routable IP Address in the host OS of the fourth instance.
D. Modify the routing table for the public subne

Answer: B

NEW QUESTION 347
You have a load balancer configured for VPC, and all back-end Amazon EC2 instances are in service. However, your web browser times out when connecting to the load baIancer's DNS name. Which options are probable causes of this behavior? Choose 2 answers

A. The load balancer was not configured to use a public sub net with an Internet gateway configured
B. The Amazon EC2 instances do not have a dynamically allocated private IP address
C. The security groups or network ACLs are not property configured for web traffic.
D. The load balancer is not configured in a private subnet with a NAT instance.
E. The VPC does not have a VGW configure

Answer: AC

NEW QUESTION 349
A company is preparing to give AWS Management Console access to developers Company policy mandates identity federation and role-based access control. Roles are currently assigned using groups in the corporate Active Directory. What combination of the following will give developers access to the AWS console? {Select 2} Choose 2 answers

A. AWS Directory Service AD Connector
B. AWS Directory Service Simple AD
C. AWS Identity and Access Management groups
D. AWS identity and Access Management roles
E. AWS identity and Access Management users

Answer: AD

NEW QUESTION 352
You are deploying an application to collect votes for a very popular television show. Millions of users will submit votes using mobile devices. The votes must be collected into a durable, scalable, and highly available data store for real-time public tabulation. Which service should you use?

A. Amazon DynamoDB
B. Amazon Redshift
C. Amazon Kinesis
D. Amazon Simple Queue Service

Answer: C

NEW QUESTION 353
You are deploying an application to track GPS coordinates of delivery trucks in the United States. Coordinates are transmitted from each delivery t ruck once every three seconds. You need to design an architecture that will enable real-time processing of these coordinates from multiple consumers. Which service should you use to implement data ingestion?

A. Amazon Kinesis
B. AWS Data Pipeline
C. Amazon AppStream
D. Amazon Simple Queue Sen/ice

Answer: A

NEW QUESTION 354
You have an application running on an Amazon Elastic Compute Cloud instance, that uploads 5 GB video objects to Amazon Simple Storage Service (53). Video uploads are taking longer than expected, resulting in poor application performance. Which method will help improve performance of your application?

A. Enable enhanced networking
B. Use Amazon 53 multipart upload
C. Leveraging Amazon CloudFront, use the HTIP POST method to reduce latency.
D. Use Amazon Elastic Block Store Provisioned IOPs and use an Amazon EBS-optimized instance

Answer:

B

NEW QUESTION 359
When will you incur costs with an Elastic IP address (EIP)?

A. When an EIP is allocated.
B. When it is allocated and associated with a running instance.
C. When it is allocated and associated with a stopped instance.
D. Costs are incurred regardless of whether the EIP is associated with a running instanc

Answer: D

NEW QUESTION 361
You try to connect via SSH to a newly created Amazon EC2 instance and get one of the following error messages:
"Network error: Connection timed out" or "Error connecting to [instance], reason: -> Connection timed out: connect,"
You have confirmed that the network and security group rules are configured correctly and the instance is passing status checks. What steps should you take to identify the source of the behavior? Choose 2 answers

A. Verify that the private key file corresponds to the Amazon EC2 key pair assigned at launch.
B. Verify that your IAM user policy has permission to launch Amazon EC2 instances.
C. Verify that you are connecting with the appropriate user name for your AMI.
D. Verify that the Amazon EC2 Instance was launched with the proper IAM role.
E. Verify that your federation trust to AWS has been establishe

Answer: AC

Explanation:
Reference:
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Troubleshooting|nstancesConnecting.html

NEW QUESTION 362
Which of the following services natively encrypts data at rest within an AWS region? Choose 2 answers

A. AWS Storage Gateway
B. Amazon DynamoDB
C. Amazon C|oudFront
D. Amazon Glacier
E. Amazon Simple Queue Senrice

Answer: AD

Explanation:
Reference:
https://media.amazonwebservices.com/AWS_Securing_Data_ at_Rest_ with_E ncryption. pdf (page 12)

NEW QUESTION 367
A company is building software on AWS that requires access to various AWS services. Which configuration should be used to ensure mat AWS credentials (i.e ., Access Key ID/Secret Access Key combination) are not compromised?

A. Enable Multi-Factor Authentication for your AWS root account.
B. Assign an IAM role to the Amazon EC2 instance.
C. Store the AWS Access Key ID/Secret Access Key combination in software comments.
D. Assign an IAM user to the Amazon EC2 Instanc

Answer: A

Explanation:
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/IAMBestPractices.htm|

NEW QUESTION 369
A t2.medium EC2 instance type must be launched with what type of Amazon Machine Image (AMI)?

A. An Instance store Hardware Virtual Machine AMI
B. An Instance store Paravirtual AMI
C. An Amazon EBS-backed Hardware Virtual Machine AMI
D. An Amazon EBS-backed Paravirtual AMI

Answer: A

Explanation:
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instanee-resize.html

NEW QUESTION 372
An existing application stores sensitive information on a non-boot Amazon EBS data volume attached to an Amazon Elastic Compute Cloud instance. Which of the following approaches would protect the sensitive data on an Amazon EBS volume?

A. Upload your customer keys to AWS CloudHS
B. Associate the Amazon EBS volume with AWS C|oudHS
C. Re-mount the Amazon EBS volume.
D. Create and mount a new, encrypted Amazon EBS volum
E. Move the data to the new volum
F. Delete the old Amazon EBS volume.
G. Unmount the EBS volum
H. Toggle the encryption attribute to Tru
I. Re-mount the Amazon EBS volume.
J. Snapshot the current Amazon EBS volum
K. Restore the snapshot to a new, encrypted Amazon EBS volum
L. Mount the Amazon EBS volume

Answer: D

NEW QUESTION 373
A US-based company is expanding their web presence into Europe. The company wants to extend their AWS infrastructure from Northern Virginia (us-east-1) into the Dublin (eu-west-1) region. Which of the following options would enable an equivalent experience for users on both continents?

A. Use a public-facing load balancer per region to load-balance web traffic, and enable HTIP health checks.
B. Use a public-facing load balancer per region to load-balance web traffic, and enable sticky sessions.
C. Use Amazon Route 53, and apply a geolocation routing policy to distribute traffic across both regions.
D. Use Amazon Route 53, and apply a weighted routing policy to distribute traffic across both region

Answer: D

Explanation:
Reference: http://docs.aws.amazon.com/Route53/latest/DeveIoperGuide/routing-policy.html

NEW QUESTION 376
Which of the following are use cases for Amazon DynamoDB? Choose 3 answers

A. Storing BLOB data.
B. Managing web sessions.
C. Storing JSON documents.
D. Storing metadata for Amazon 53 objects.
E. Running relational joins and complex updates.
F. Storing large amounts of infrequently accessed dat

Answer: CEF

NEW QUESTION 380
......

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