

# Amazon-Web-Services

## Exam Questions SAA-C01

AWS Certified Solutions Architect - Associate



### NEW QUESTION 1

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 IP multicast The database tier uses shared-storage clustering to provide database fail over capability, and uses several read slaves for scaling Data on all servers 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 servers: store read-only data in S3, and copy from S3 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 servers: 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 multiclas
- G. Database: use RDS with multi-AZ deployment and one or more Read Replica
- H. Backup: web and app servers backed up weekly via AMIs, database backed up via DB snapshots.
- I. Web servers: store read-only data in S3, and copy from S3 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 via AMIs, database backed up via DB snapshots.
- M. Web servers: store read-only data in S3, and copy from S3 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 servers backed up weekly via AMIs, database backed up via DB snapshots.

**Answer: A**

#### Explanation:

<https://d0.awsstatic.com/whitepapers/Storage/AWS%20Storage%20Services%20Whitepaper-v9.pdf> Amazon Glacier doesn't suit all storage situations. Listed following are a few storage needs for which you should consider other AWS storage options instead of Amazon Glacier.

Data that must be updated very frequently might be better served by a storage solution with lower read/write latencies, such as Amazon EBS, Amazon RDS, Amazon DynamoDB, or relational databases running on EC2.

### NEW QUESTION 2

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 individual 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 AMIs and supplement with file-level backup to S3 using traditional enterprise backup software to provide file level restore
- B. Backup RDS using a Multi-AZ Deployment Backup the EC2 instances using Amis, and supplement by copying file system data to S3 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 S3 using Oracle RMAN Backup the EC2 instances using Amis, and supplement with EBS snapshots for individual volume restore.

**Answer: A**

#### Explanation:

You need to use enterprise backup software to provide file level restore. See

[https://d0.awsstatic.com/whitepapers/Backup\\_and\\_Recovery\\_Approaches\\_Using\\_AWS.pdf](https://d0.awsstatic.com/whitepapers/Backup_and_Recovery_Approaches_Using_AWS.pdf) Page 18:

If your existing backup software does not natively support the AWS cloud, you can use AWS storage gateway products. AWS Storage Gateway is a virtual appliance that provides seamless and secure integration between your data center and the AWS storage infrastructure.

### NEW QUESTION 3

You need a persistent and durable storage to trace call activity of an IVR (Interactive Voice Response) system. Call duration is mostly in the 2-3 minutes timeframe. Each traced call can be either active or terminated. An external application needs to know each minute the list of currently active calls, which are usually a few calls/second. Put once per month there is a periodic peak up to 1000 calls/second for a few hours. The system is open 24/7 and any downtime should be avoided. Historical data is periodically archived to files. Cost saving is a priority for this project. What database implementation would better fit this scenario, keeping costs as low as possible?

- A. Use RDS Multi-AZ with two tables, one for "-Active calls" and one for "-Terminated calls". In this way the "Active calls\_ table is always small and effective to access.
- B. Use DynamoDB with a "Calls" table and a Global Secondary Index on a "IsActive" attribute that is present for active calls only In this way the Global Secondary index is sparse and more effective.
- C. Use DynamoDB with a "Calls" table and a Global secondary index on a "State" attribute that can equal to "active" or "terminated" in this way the Global Secondary index can be used for all Items in the table.
- D. Use RDS Multi-AZ with a "CALLS" table and an Indexed "STATE\* field that can be equal to 'ACTIVE" or "-TERMINATED" In this way the SOL query Is optimized by the use of the Index.

**Answer: B**

#### Explanation:

Q: Can a global secondary index key be defined on non-unique attributes?

Yes. Unlike the primary key on a table, a GSI index does not require the indexed attributes to be unique.

Q: Are GSI key attributes required in all items of a DynamoDB table?

No. GSIs are sparse indexes. Unlike the requirement of having a primary key, an item in a DynamoDB table does not have to contain any of the GSI keys. If a GSI key has both hash and range elements, and a table item omits either of them, then that item will not be indexed by the corresponding GSI.

In such cases, a GSI can be very useful in efficiently locating items that have an uncommon attribute.

#### NEW QUESTION 4

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

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

**Answer:** AB

#### Explanation:

The Route 53 and IAM are global.

As per the document defined, new IPs should be reserved not the same ones. Elastic IP Addresses are static IP addresses designed for dynamic cloud computing. Unlike traditional static IP addresses, however, Elastic IP addresses enable you to mask instance or Availability Zone failures by programmatically remapping your public IP addresses to instances in your account in a particular region. For DR, you can also pre-allocate some IP addresses for the most critical systems so that their

IP addresses are already known before disaster strikes. This can simplify the execution of the DR plan.

#### NEW QUESTION 5

An International company has deployed a multi-tier web application that relies on DynamoDB in a single region. For regulatory reasons they need disaster recovery capability in a separate region with a Recovery Time Objective of 2 hours and a Recovery Point Objective of 24 hours. They should synchronize their data on a regular basis and be able to provision the web application rapidly using CloudFormation.

The objective is to minimize changes to the existing web application, control the throughput of DynamoDB used for the synchronization of data and synchronize only the modified elements. Which design would you choose to meet these requirements?

- A. Use AWS data Pipeline to schedule a DynamoDB cross region copy once a day
- B. create a 'LastUpdated' attribute in your DynamoDB table that would represent the timestamp of the last update and use it as a filter.
- C. Use EMR and write a custom script to retrieve data from DynamoDB in the current region using a SCAN operation and push it to DynamoDB in the second region.
- D. Use AWS data Pipeline to schedule an export of the DynamoDB table to S3 in the current region once a day then schedule another task immediately after it that will import data from S3 to DynamoDB in the other region.
- E. Send also each item into an SQS queue in the second region; use an auto-scaling group behind the SQS queue to replay the write in the second region.

**Answer:** A

#### NEW QUESTION 6

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

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

**Answer:** C

#### Explanation:

Q. Can I use AWS Direct Connect and a VPN Connection to the same VPC simultaneously?

Yes. However, only in fail-over scenarios. The Direct Connect path will always be preferred, when established, regardless of AS path prepending.

<https://aws.amazon.com/directconnect/faqs/>

#### NEW QUESTION 7

You are designing Internet connectivity for your VPC. The Web servers must be available on the Internet. The application must have a highly available architecture.

Which alternatives should you consider? (Choose two.)

- A. Configure a NAT instance in your VPC
- B. Create a default route via the NAT instance and associate it with all subnets
- C. Configure a DNS A record that points to the NAT instance public IP address.
- D. Configure a CloudFront distribution and configure the origin to point to the private IP addresses of your Web server
- E. Configure a Route53 CNAME record to your CloudFront distribution.
- F. Place all your web servers behind ELB. Configure a Route53 CNAME to point to the ELB DNS name.
- G. Assign EIPs to all web servers
- H. Configure a Route53 record set with all EIPs with health checks and DNS failover.
- I. Configure ELB with an EIP
- J. Place all your Web servers behind ELB
- K. Configure a Route53 A record that points to the EIP.

**Answer:** CD

### NEW QUESTION 8

A corporate web application is deployed within an Amazon Virtual Private Cloud (VPC) and is connected to the corporate data center via an IPsec VPN. The application must authenticate against the on-premises LDAP server. After authentication, each logged-in user can only access an Amazon Simple Storage Space (S3) keyspace specific to that user.

Which two approaches can satisfy these objectives? (Choose two.)

- A. Develop an identity broker that authenticates against IAM security Token service to assume a IAM role in order to get temporary AWS security credentials The application calls the identity broker to get AWS temporary security credentials with access to the appropriate S3 bucket.
- B. The application authenticates against LDAP and retrieves the name of an IAM role associated with the user.
- C. The application then calls the IAM Security Token Service to assume that IAM role.
- D. The application can use the temporary credentials to access the appropriate S3 bucket.
- E. Develop an identity broker that authenticates against LDAP and then calls IAM Security Token Service to get IAM federated user credentials.
- F. The application calls the identity broker to get IAM federated user credentials with access to the appropriate S3 bucket.
- G. The application authenticates against LDAP the application, then calls the AWS identity and Access Management (IAM) Security service to log in to IAM using the LDAP credentials, the application can use the IAM temporary credentials to access the appropriate S3 bucket.
- H. The application authenticates against IAM Security Token Service using the LDAP credentials, the application uses those temporary AWS security credentials to access the appropriate S3 bucket.

**Answer: BC**

#### Explanation:

Imagine that in your organization, you want to provide a way for users to copy data from their computers to a backup folder. You build an application that users can run on their computers. On the back end, the application reads and writes objects in an S3 bucket. Users don't have direct access to AWS. Instead, the application communicates with an identity provider (IdP) to authenticate the user. The IdP gets the user information from your organization's identity store (such as an LDAP directory) and then generates a SAML assertion that includes authentication and authorization information about that user. The application then uses that assertion to make a call to the AssumeRoleWithSAML API to get temporary security credentials. The app can then use those credentials to access a folder in the S3 bucket that's specific to the user. [http://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_roles\\_providers\\_saml.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_saml.html)

### NEW QUESTION 9

Your company has an on-premises multi-tier PHP web application, which recently experienced downtime due to a large burst in web traffic due to a company announcement. Over the coming days, you are expecting similar announcements to drive similar unpredictable bursts, and are looking to find ways to quickly improve your infrastructure's ability to handle unexpected increases in traffic. The application currently consists of 2 tiers a web tier which consists of a load balancer and several Linux Apache web servers as well as a database tier which hosts a Linux server hosting a MySQL database.

Which scenario below will provide full site functionality, while helping to improve the ability of your application in the short timeframe required?

- A. Failover environment: Create an S3 bucket and configure it for website hosting.
- B. Migrate your DNS to Route53 using zone file import, and leverage Route53 DNS failover to failover to the S3 hosted website.
- C. Hybrid environment: Create an AMI, which can be used to launch web servers in EC2. Create an Auto Scaling group, which uses the AMI to scale the web tier based on incoming traffic.
- D. Leverage Elastic Load Balancing to balance traffic between on-premises web servers and those hosted in AWS.
- E. Offload traffic from on-premises environment: Setup a CloudFront distribution, and configure CloudFront to cache objects from a custom origin.
- F. Choose to customize your object cache behavior, and select a TTL that objects should exist in cache.
- G. Migrate to AWS: Use VM Import/Export to quickly convert an on-premises web server to an AMI.
- H. Create an Auto Scaling group, which uses the imported AMI to scale the web tier based on incoming traffic.
- I. Create an RDS read replica and setup replication between the RDS instance and on-premises MySQL server to migrate the database.

**Answer: C**

#### Explanation:

You can have CloudFront sit in front of your on-prem web environment, via a custom origin (the origin doesn't have to be in AWS). This would protect against unexpected bursts in traffic by letting CloudFront handle the traffic that it can out of cache, thus hopefully removing some of the load from your on-prem web servers.

### NEW QUESTION 10

You're running an application on-premises due to its dependency on non-x86 hardware and want to use AWS for data backup. Your backup application is only able to write to POSIX-compatible block-based storage. You have 140TB of data and would like to mount it as a single folder on your file server.

Users must be able to access portions of this data while the backups are taking place. What backup solution would be most appropriate for this use case?

- A. Use Storage Gateway and configure it to use Gateway Cached volumes.
- B. Configure your backup software to use S3 as the target for your data backups.
- C. Configure your backup software to use Glacier as the target for your data backups.
- D. Use Storage Gateway and configure it to use Gateway Stored volume.

**Answer: D**

#### Explanation:

Data is hosted on the On-premise server as well. The requirement for 140TB is for file server On-Premise more to confuse and not in AWS. Just need a backup solution hence stored instead of cached volumes.

### NEW QUESTION 10

Your department creates regular analytics reports from your company's log files. All log data is collected in Amazon S3 and processed by daily Amazon Elastic MapReduce (EMR) jobs that generate daily PDF reports and aggregated tables in CSV format for an Amazon Redshift data warehouse. Your CFO requests that you optimize the cost structure for this system.

Which of the following alternatives will lower costs without compromising average performance of the system or data integrity for the raw data?

- A. Use reduced redundancy storage (RRS) for all data in S3. Use a combination of Spot Instances and Reserved Instances for Amazon EMR job.
- B. Use Reserved Instances for Amazon Redshift.
- C. Use reduced redundancy storage (RRS) for PDF and .csv data in S3. Add Spot Instances to EMR job.
- D. Use Spot Instances for Amazon Redshift.

- E. Use reduced redundancy storage (RRS) for PDF and .csv data In Amazon S3. Add Spot Instances to Amazon EMR job
- F. Use Reserved Instances for Amazon Redshift.
- G. Use reduced redundancy storage (RRS) for all data in Amazon S3. Add Spot Instances to Amazon EMR job
- H. Use Reserved Instances for Amazon Redshift.

**Answer: D**

**Explanation:**

Reserved Instances (a.k.a. Reserved Nodes) are appropriate for steady-state production workloads, and offer significant discounts over On-Demand pricing.  
<https://aws.amazon.com/redshift>

Q: What are some EMR best practices?

If you are running EMR in production you should specify an AMI version, Hive version, Pig version, etc. to make sure the version does not change unexpectedly (e.g. when EMR later adds support for a newer version). If your cluster is mission critical, only use Spot instances for task nodes because if the Spot price increases you may lose the instances. In development, use logging and enable debugging to spot and correct errors faster. If you are using GZIP, keep your file size to 1–2 GB because GZIP files cannot be split. Click here to download the white paper on Amazon EMR best practices. <https://aws.amazon.com/elasticmapreduce/faqs>

**NEW QUESTION 15**

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

- A. Use OAuth 2.0 to retrieve temporary AWS security credentials to enable your NOC members to sign in to the AVVS Management Console.
- B. Use web Identity Federation to retrieve AWS temporary security credentials to enable your NOC members to sign in to the AWS Management Console.
- C. Use your on-premises SAML 2.0-compliant identity provider (IDP) to grant the NOC members federated access to the AWS Management Console via the AWS single sign-on (SSO) endpoint.
- D. Use your on-premises SAML2.0-compliant identity provider (IDP) to retrieve temporary security credentials to enable NOC members to sign in to the AWS Management Console.

**Answer: C**

**NEW QUESTION 17**

You are designing an SSUTLS solution that requires HTTPS clients to be authenticated by the Web server using client certificate authentication. The solution must be resilient.

Which of the following options would you consider for configuring the web server infrastructure? (Choose two.)

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

**Answer: AB**

**Explanation:**

This question is regarding “two-way” SSL authentication.

Currently, ELBs cannot support authentication for the client side SSL/TLS cert required for two-way SSL authentication to succeed. Therefore, you only have two options:

\A. Configure the ELB with a TCP/443 listener. This is effectively TLS “pass through” mode, where the TLS connection does not terminate on the ELB, it is passed through and decrypted on the back-end servers. This will cause quite a bit of CPU overhead on the back-end instances, due to the lack of TLS offload that cannot happen on the ELB, so an auto-scaling group which monitors the web server CPU metrics would be essential here. (Not that you probably wouldn't have it anyway, just saying!)

\B. Don't use an ELB. Just have the web servers act as the endpoint for the traffic, and let Route53 DNS serve in the place of the ELB by load balancing client DNS queries across the web servers. C and D are not options here, since neither are supported by AWS.

**NEW QUESTION 18**

You are designing a social media site and are considering how to mitigate distributed denial-of-service (DDoS) attacks. Which of the below are viable mitigation techniques? (Choose three.)

- A. Add multiple elastic network interfaces (ENIs) to each EC2 instance to increase the network bandwidth.
- B. Use dedicated instances to ensure that each instance has the maximum performance possible.
- C. Use an Amazon CloudFront distribution for both static and dynamic content.
- D. Use an Elastic Load Balancer with auto scaling groups at the we
- E. App and Amazon Relational Database Service (RDS) tiers
- F. Add alert Amazon CloudWatch to look for high Network in and CPU utilization.
- G. Create processes and capabilities to quickly add and remove rules to the instance OS firewall

**Answer: CDE**

**NEW QUESTION 20**

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

- A. Stateless instances for the web and application tier synchronized using ElastiCache Memcached in an autoscaimg group monitored with CloudWatc
- B. And RDSwith read replicas.
- C. Stateful instances for the web and application tier in an autoscaling group monitored with CloudWatch and RDS with read replicas.
- D. Stateful instances for the web and application tier in an autoscaling group monitored with CloudWatc
- E. And multi-AZ RDS.
- F. Stateless instances for the web and application tier synchronized using ElastiCache Memcached in an autoscaling group monitored with CloudWatch and multi-

AZ RDS.

**Answer:** A

**Explanation:**

“A readonly reporting site” - so stateless and read-replicas can be used to scale. Multi-AZ will not provide the scaling requirements.

**NEW QUESTION 21**

You are running a news website in the eu-west-1 region that updates every 15 minutes. The website has a world-wide audience it uses an Auto Scaling group behind an Elastic Load Balancer and an Amazon RDS database. Static content resides on Amazon S3, and is distributed through Amazon CloudFront. Your Auto Scaling group is set to trigger a scale up event at 60% CPU utilization, you use an Amazon RDS extra large DB instance with 10,000 Provisioned IOPS its CPU utilization is around 80%. While freeable memory is in the 2 GB range.

Web analytics reports show that the average load time of your web pages is around 1.5 to 2 seconds, but your SEO consultant wants to bring down the average load time to under 0.5 seconds. How would you improve page load times for your users? (Choose three.)

- A. Lower the scale up trigger of your Auto Scaling group to 30% so it scales more aggressively.
- B. Add an Amazon ElastiCache caching layer to your application for storing sessions and frequent DB queries
- C. Configure Amazon CloudFront dynamic content support to enable caching of re-usable content from your site
- D. Switch Amazon RDS database to the high memory extra large Instance type
- E. Set up a second installation in another region, and use the Amazon Route 53 latency-based routing feature to select the right region.

**Answer:** BCE

**Explanation:**

The freeable memory includes the amount of physical memory left unused by the system plus the total amount of buffer or page cache memory that are free and available.

So it's freeable memory across the entire system. While MySQL is the main consumer of memory on the host we do have internal processes in addition to the OS that use up a small amount of additional memory.

If you see your freeable memory near 0 or also start seeing swap usage then you may need to scale up to a larger instance class or adjust MySQL memory settings. For example decreasing

the `innodb_buffer_pool_size` (by default set to 75% of physical memory) is one way example of adjusting MySQL memory settings

Takeaway: extra mem is not going to help page load times here, but a 2nd region might. Keep in mind they're going for a 66%-75% reduction in page load times – what if you added a region in Australia or HK, would that not help your worldwide users? rather than having traffic go to us-east.

**NEW QUESTION 23**

A large real-estate brokerage is exploring the option of adding a cost-effective location based alert to their existing mobile application. The application backend infrastructure currently runs on AWS. Users who opt in to this service will receive alerts on their mobile device regarding real-estate offers in proximity to their location. For the alerts to be relevant delivery time needs to be in the low minute count the existing mobile app has 5 million users across the us. Which one of the following architectural suggestions would you make to the customer?

- A. The mobile application will submit its location to a web service endpoint utilizing Elastic Load Balancing and EC2 instances: DynamoDB will be used to store and retrieve relevant offers. EC2 instances will communicate with mobile carriers/device providers to push alerts back to mobile application.
- B. Use AWS DirectConnect or VPN to establish connectivity with mobile carriers. EC2 instances will receive the mobile applications' location through carrier connection: ROS will be used to store and relevant offers. EC2 instances will communicate with mobile carriers to push alerts back to the mobile application.
- C. The mobile application will send device location using SQS.
- D. EC2 instances will retrieve the relevant offers from DynamoDB. AWS Mobile Push will be used to send offers to the mobile application.
- E. The mobile application will send device location using AWS Mobile Push. EC2 instances will retrieve the relevant offers from DynamoDB. EC2 instances will communicate with mobile carriers/device providers to push alerts back to the mobile application.

**Answer:** A

**Explanation:**

AWS using SQS to store the message from mobile apps, and using AWS Mobile Push to send offers to mobile apps.

**NEW QUESTION 28**

A company is building a voting system for a popular TV show, viewers watch the performances then visit the show's website to vote for their favorite performer. It is expected that in a short period of time after the show has finished the site will receive millions of visitors. The visitors will first login to the site using their Amazon.com credentials and then submit their vote. After the voting is completed the page will display the vote totals. The company needs to build the site such that can handle the rapid influx of traffic while maintaining good performance but also wants to keep costs to a minimum. Which of the design patterns below should they use?

- A. Use CloudFront and an Elastic Load balancer in front of an auto-scaled set of web servers, the web servers will first call the Login With Amazon service to authenticate the user then process the users vote and store the result into a multi-AZ Relational Database Service instance.
- B. Use CloudFront and the static website hosting feature of S3 with the Javascript SDK to call the Login With Amazon service to authenticate the user, use IAM Roles to gain permissions to a DynamoDB table to store the users vote.
- C. Use CloudFront and an Elastic Load Balancer in front of an auto-scaled set of web servers, the web servers will first call the Login with Amazon service to authenticate the user, the web servers will process the users vote and store the result into a DynamoDB table using IAM Roles for EC2 instances to gain permissions to the DynamoDB table.
- D. Use CloudFront and an Elastic Load Balancer in front of an auto-scaled set of web servers, the web servers will first call the Login With Amazon service to authenticate the user, the web servers will process the users vote and store the result into an SQS queue using IAM Roles for EC2 Instances to gain permissions to the SQS queue.
- E. With Amazon service to authenticate the user, the web servers will process the users vote and store the result into a DynamoDB table.
- F. A set of application servers will then retrieve the items from the queue and store the result into a DynamoDB table.

**Answer:** D

**NEW QUESTION 32**

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 servers 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 S3 use EMR to apply heuristics on the logs

**Answer: B**

**Explanation:**

Amazon Kinesis Streams allows for real-time data processing. With Amazon Kinesis Streams, you can continuously collect data as it is generated and promptly react to critical information about your business and operations.

<https://aws.amazon.com/kinesis/streams/>

**NEW QUESTION 33**

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 IP Address

**Answer: C**

**NEW QUESTION 35**

What does Amazon SWF stand for?

- A. Simple Web Flow
- B. Simple Work Flow
- C. Simple Wireless Forms
- D. Simple Web Form

**Answer: B**

**NEW QUESTION 39**

Fill in the blanks: Resources that are created in AWS are identified by a unique identifier called an \_\_\_\_\_

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

**Answer: C**

**NEW QUESTION 44**

When you view the block device mapping for your instance, you can see only the EBS volumes, not the instance store volumes.

- A. Depends on the instance type
- B. FALSE
- C. Depends on whether you use API call
- D. TRUE

**Answer: D**

**Explanation:**

When you view the block device mapping for your instance, you can see only the EBS volumes, not the instance store volumes. You can use instance metadata to query the complete block device mapping. The base URI for all requests for instance metadata is

<http://169.254.169.254/latest/>. <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/block-device-mappingconcepts.html#bdm-instance-metadata>

**NEW QUESTION 46**

What are the initial settings of a user created security group?

- A. Allow all inbound traffic and Allow no outbound traffic
- B. Allow no inbound traffic and Allow no outbound traffic
- C. Allow no inbound traffic and Allow all outbound traffic
- D. Allow all inbound traffic and Allow all outbound traffic

**Answer: C**

**NEW QUESTION 48**

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

- A. TRUE
- B. FALSE

**Answer:** A

**Explanation:**

**AWS Direct Connect Keys**

AWS Direct Connect implements the following policy keys:

- `aws:CurrentTime` (for date/time conditions)
- `aws:EpochTime` (the date in epoch or UNIX time, for use with date/time conditions)
- `aws:SecureTransport` (Boolean representing whether the request was sent using SSL)
- `aws:SourceIp` (the requester's IP address, for use with IP address conditions)
- `aws:UserAgent` (information about the requester's client application, for use with string conditions)

If you use `aws:SourceIp`, and the request comes from an Amazon EC2 instance, the instance's public IP address is used to determine if access is allowed.

**Note**

For services that use only SSL, such as Amazon Relational Database Service and Amazon Route 53, the `aws:SecureTransport` key has no meaning.

Key names are case-insensitive. For example, `aws:CurrentTime` is equivalent to `AWS:currenttime`.

[http://docs.aws.amazon.com/directconnect/latest/UserGuide/using\\_iam.html](http://docs.aws.amazon.com/directconnect/latest/UserGuide/using_iam.html)

**NEW QUESTION 49**

Can we attach an EBS volume to more than one EC2 instance at the same time?

- A. No
- B. Yes.
- C. Only EC2-optimized EBS volumes.
- D. Only in read mod

**Answer:** A

**NEW QUESTION 54**

Fill in the blanks: The base URI for all requests for instance metadata is

- A. <http://254.169.169.254/latest/>
- B. <http://169.169.254.254/latest/>
- C. <http://127.0.0.1/latest/>
- D. <http://169.254.169.254/latest/>

**Answer:** D

**Explanation:**

<http://aws.amazon.com/search?searchQuery=metadata&searchPath=all&x=0&y=0>

**NEW QUESTION 59**

Amazon RDS DB snapshots and automated backups are stored in

- A. Amazon S3
- B. Amazon ECS Volume
- C. Amazon RDS
- D. Amazon EMR

**Answer:** A

**NEW QUESTION 61**

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

**Explanation:**

Groups can't be nested; they can contain only users, not other groups. [http://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_groups.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/id_groups.html)

**NEW QUESTION 64**

Changes to the backup window take effect \_\_\_\_\_.

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

**Answer:** C

**Explanation:**

Changes to the backup window take effect immediately, with the limitations that the specified backup window must be at least 10 minutes in the future, and the backup window cannot overlap with the weekly maintenance window for the instance.

**NEW QUESTION 69**

Using Amazon CloudWatch's Free Tier, what is the frequency of metric updates which you receive?

- A. 5 minutes
- B. 500 milliseconds.
- C. 30 seconds
- D. 1 minute

**Answer:** A

**Explanation:**

You can get started with Amazon CloudWatch for free. Many applications should be able to operate within these free tier limits.

New and existing customers also receive 3 dashboards of up to 50 metrics each per month at no additional charge. Basic Monitoring metrics (at five-minute frequency) for Amazon EC2 instances are free of charge, as are all metrics for Amazon EBS volumes, Elastic Load Balancers, and Amazon RDS DB instances.

<https://aws.amazon.com/cloudwatch/pricing/>

**NEW QUESTION 72**

Which is the default region in AWS?

- A. eu-west-1
- B. us-east-1
- C. us-east-2
- D. ap-southeast-1

**Answer:** B

**Explanation:**

All the main AWS services (except Route 53 & CloudFront) allow you to select which region you would like to use. The US East (N. Virginia) is the default region. You can change the region by using the dropdown menu in the top right of the management console.

**NEW QUESTION 74**

EBS Snapshots occur \_\_\_\_\_

- A. Asynchronously
- B. Synchronously
- C. Weekly

**Answer:** A

**Explanation:**

Snapshots occur asynchronously; the point-in-time snapshot is created immediately, but the status of the snapshot is pending until the snapshot is complete (when all of the modified blocks have been transferred to Amazon S3), which can take several hours for large initial snapshots or subsequent snapshots where many blocks have changed. <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-creating-snapshot.html>

**NEW QUESTION 76**

Disabling automated backups disable the point-in-time recovery.

- A. if configured to can
- B. will never
- C. will

**Answer:** C

**NEW QUESTION 81**

How can I change the security group membership for interfaces owned by other AWS, such as Elastic Load Balancing?

- A. By using the service specific console or API/CLI commands
- B. None of these
- C. Using Amazon EC2 API/CLI
- D. using all these methods

**Answer:** A

**Explanation:**

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/elb-security-groups.html>

### Security Groups for Load Balancers in a VPC

When you use the AWS Management Console to create a load balancer in a VPC, you can choose an existing security group for the VPC or create a new security group for the VPC. If you choose an existing security group, it must allow traffic in both directions to the listener and health check ports for the load balancer. If you choose to create a security group, the console automatically adds rules to allow all traffic on these ports.

[Nondefault VPC] If you use the **AWS CLI or API** to create a load balancer in a nondefault VPC, but you don't specify a security group, your load balancer is automatically associated with the default security group for the VPC.

[Default VPC] If you use the **AWS CLI or API** to create a load balancer in your default VPC, you can't choose an existing security group for your load balancer. Instead, Elastic Load Balancing provides a security group with rules to allow all traffic on the ports specified for the load balancer. Elastic Load Balancing creates only one such security group per AWS account, with a name of the form `default_elb_id` (for example, `default_elb_fc5fbed3-0405-3b7d-a328-aa290EXAMPLE`). Subsequent load balancers that you create in the default VPC also use this security group. Be sure to review the security group rules to ensure that they allow traffic on the listener and health check ports for the new load balancer. When you delete your load balancer, this security group is not deleted automatically.

If you add a listener to an existing load balancer, you must review your security groups to ensure they allow traffic on the new listener port in both directions.

#### NEW QUESTION 85

Can a 'user' be associated with multiple AWS accounts?

- A. No
- B. Yes

**Answer: A**

#### NEW QUESTION 88

What is the durability of S3 RRS?

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

**Answer: A**

#### Explanation:

RRS = Reduced Redundancy Storage

	Standard	Standard - Infrequent Access	Reduced Redundancy Storage
Durability	99.999999999%	99.999999999%	99.99%

#### NEW QUESTION 93

What does specifying the mapping `/dev/sdc=none` when launching an instance do?

- A. Prevents `/dev/sdc` from creating the instance.
- B. Prevents `/dev/sdc` from deleting the instance.
- C. Set the value of `/dev/sdc` to 'zero'.
- D. Prevents `/dev/sdc` from attaching to the instance.

**Answer: D**

#### Explanation:

`none` - Suppresses an existing mapping of the device from the AMI used to launch the instance. For example: `"/dev/sdc=none"`.  
<http://docs.aws.amazon.com/AWSEC2/latest/CommandLineReference/ApiReference-cmd-RegisterImage.html>

#### NEW QUESTION 98

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

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

**Answer: D**

#### Explanation:

DB Instance detail - You just enable that your DB instance can be deployed in Multi-AZ. However, you select the availability zone (Which AZ will be for primary and which one will be for secondary) in Additional configuration.

**NEW QUESTION 102**

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

- A. Under DB INSTANCE DETAILS
- B. Under REVIEW
- C. Under MANAGEMENT OPTIONS
- D. Under ENGINE SELECTION

**Answer: C**

**NEW QUESTION 104**

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**

**Explanation:**

AWS managed policies are designed to provide permissions for many common use cases. For example, there are AWS managed policies that define typical permissions for administrators (all access), for power users (all access except IAM), and for other various levels of access to AWS services. AWS managed policies make it easier for you to assign appropriate permissions to users, groups, and roles than if you had to write the policies yourself.  
[http://docs.aws.amazon.com/IAM/latest/UserGuide/access\\_policies\\_managed-vs-inline.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_managed-vs-inline.html)

**NEW QUESTION 105**

Does Amazon RDS allow direct host access via Telnet, Secure Shell (SSH), or Windows Remote Desktop Connection?

- A. Yes
- B. No
- C. Depends on if it is in VPC or not

**Answer: B**

**Explanation:**

In order to deliver a managed service experience, Amazon RDS does not provide shell access to DB instances, and it restricts access to certain system procedures and tables that require advanced privileges. Amazon RDS supports access to databases on a DB instance using any standard SQL client application such as Microsoft SQL Server Management Studio. **Amazon RDS does not allow direct host access to a DB instance via Telnet, Secure Shell (SSH), or Windows Remote Desktop Connection.** When you create a DB instance, you are assigned to the `db_owner` role for all databases on that instance, and you will have all database-level permissions except for those that are used for backups (Amazon RDS manages backups for you).

**NEW QUESTION 109**

To view information about an Amazon EBS volume, open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>, click \_\_\_\_\_ in the Navigation pane.

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

**Answer: D**

**Explanation:**

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-describing-volumes.html>

**Viewing Volume Information**

You can view descriptive information for your Amazon EBS volumes in a selected region at a time in the AWS Management Console. You can also view detailed information about a single volume, including the size, volume type, whether or not the volume is encrypted, which master key was used to encrypt the volume, and the specific instance to which the volume is attached.

**To view information about an EBS volume using the console**

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, choose **Volumes**.
3. To view more information about a volume, select it.

**NEW QUESTION 112**

Using Amazon IAM, can I give permission based on organizational groups?

- A. Yes but only in certain cases

- B. No
- C. Yes always

**Answer:** C

**Explanation:**

An IAM group is a collection of IAM users. You can use groups to specify permissions for a collection of users, which can make those permissions easier to manage for those users. <http://docs.aws.amazon.com/IAM/latest/UserGuide/id.html>

**NEW QUESTION 117**

What is an isolated database environment running in the cloud (Amazon RDS) called?

- A. DB Instance
- B. DB Server
- C. DB Unit
- D. DB Volume

**Answer:** A

**NEW QUESTION 119**

In the 'Detailed' monitoring data available for your Amazon EBS volumes, Provisioned IOPS volumes automatically send \_\_\_\_\_ minute metrics to Amazon CloudWatch.

- A. 3
- B. 1
- C. 5
- D. 2

**Answer:** B

**NEW QUESTION 120**

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 charge
- C. It is free.
- D. USD 0.02 per GB
- E. USD 0.01 per GB

**Answer:** B

**Explanation:**

For data transferred between an Amazon EC2 instance and Amazon RDS DB Instance in different Availability Zones of the same Region, there is no Data Transfer charge for traffic in or out of the Amazon RDS DB Instance. References:

**NEW QUESTION 125**

Are Reserved Instances available for Multi-AZ Deployments?

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

**Answer:** B

**Explanation:**

<https://aws.amazon.com/rds/faqs/>

**NEW QUESTION 127**

Which Amazon Storage behaves like raw, unformatted, external block devices that you can attach to your instances?

- A. None of these.
- B. Amazon Instance Storage
- C. Amazon EBS
- D. All of these

**Answer:** C

**NEW QUESTION 131**

MySQL installations default to port \_\_\_\_\_ .

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

**Answer:** A

**Explanation:**

[http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER\\_ConnectToInstance.html](http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ConnectToInstance.html)

**NEW QUESTION 135**

If I want to run a database in an Amazon instance, which is the most recommended Amazon storage option?

- A. Amazon Instance Storage
- B. Amazon EBS
- C. You can't run a database inside an Amazon instance.
- D. Amazon S3

**Answer:** B

**Explanation:**

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Storage.html>

**NEW QUESTION 139**

True or False: If you add a tag that has the same key as an existing tag on a DB Instance, the new value overwrites the old value.

- A. FALSE
- B. TRUE

**Answer:** B

**Explanation:**

[http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using\\_Tags.html](http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html)

**NEW QUESTION 144**

Can I use Provisioned IOPS with VPC?

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

**Answer:** D

**NEW QUESTION 149**

Fill in the blanks: "To ensure failover capabilities, consider using a \_\_\_\_ for incoming traffic on a network interface".

- A. primary public IP
- B. secondary private IP
- C. secondary public IP
- D. add on secondary IP

**Answer:** B

**Explanation:**

To ensure failover capabilities, consider using a secondary private IP for incoming traffic on an elastic network interface. In the event of an instance failure, you can move the interface and/or secondary private IP address to a standby instance

**NEW QUESTION 154**

What can I access by visiting the URL: <http://status.aws.amazon.com/>?

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

**Answer:** C

**NEW QUESTION 157**

Please select the Amazon EC2 resource which cannot be tagged.

- A. images (AMIs, kernels, RAM disks)
- B. Amazon EBS volumes
- C. Elastic IP addresses
- D. VPCs

**Answer:** C

**Explanation:**

[http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using\\_Tags.html#tag-restrictions](http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html#tag-restrictions)

Resource	Tagging support	Tagging restrictions
AMI	Yes	None
Bundle task	No	
Customer gateway	Yes	None
Dedicated Host	No	
DHCP option	Yes	None
EBS volume	Yes	None
Instance store volume	No	
Elastic IP	No	
Egress-only Internet gateway	No	
Instance	Yes	None
Internet gateway	Yes	None
Key pair	No	
NAT gateway	No	
Network ACL	Yes	None
Network interface	Yes	None
Placement group	No	
Reserved Instance	Yes	None

**NEW QUESTION 158**

How are the EBS snapshots saved on Amazon S3?

- A. Exponentially
- B. Incrementally
- C. EBS snapshots are not stored in the Amazon S3
- D. Decrementally

**Answer: B**

**NEW QUESTION 162**

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

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

**Answer: A**

**Explanation:**

Monitoring Volumes with CloudWatch

CloudWatch metrics are statistical data that you can use to view, analyze, and set alarms on the operational behavior of your volumes. The following table describes the types of monitoring data available for your Amazon EBS volumes:

**Basic**  
 Data is available automatically in 5-minute periods at no charge. This includes data for the root device volumes for EBS- backed instances.

**Detailed**  
 Provisioned IOPS SSD (io1) volumes automatically send one-minute metrics to CloudWatch.  
<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/monitoring-volume-status.html>

**Monitoring Volumes with CloudWatch**

CloudWatch metrics are statistical data that you can use to view, analyze, and set alarms on the operational behavior of your volumes.

The following table describes the types of monitoring data available for your Amazon EBS volumes.

Type	Description
Basic	Data is available automatically in 5-minute periods at no charge. This includes data for the root device volumes for EBS-backed instances.
Detailed	Provisioned IOPS SSD (io1) volumes automatically send one-minute metrics to CloudWatch.

**NEW QUESTION 163**

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

- A. 120 seconds
- B. 1 hour
- C. 10 minutes
- D. 12 hours

**Answer: B**

**Explanation:**

<https://aws.amazon.com/premiumsupport/features/>

**NEW QUESTION 164**

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**

**Explanation:**

[https://aws.amazon.com/documentation/iam/?nc1=h\\_ls](https://aws.amazon.com/documentation/iam/?nc1=h_ls)

**NEW QUESTION 165**

True or False: Without IAM, you cannot control the tasks a particular user or system can do and what AWS resources they might use.

- A. FALSE
- B. TRUE

**Answer: B**

**Explanation:**

<http://docs.aws.amazon.com/IAM/latest/UserGuide/getting-setup.html>

**NEW QUESTION 166**

Select the correct set of options. These are the initial settings for the default security group:

- A. Allow no inbound traffic, Allow all outbound traffic and Allow instances associated with this security group to talk to each other
- B. Allow all inbound traffic, Allow no outbound traffic and Allow instances associated with this security group to talk to each other
- C. Allow no inbound traffic, Allow all outbound traffic and Does NOT allow instances associated with this security group to talk to each other
- D. Allow all inbound traffic, Allow all outbound traffic and Does NOT allow instances associated with this security group to talk to each other

**Answer: A**

**Explanation:**

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html#defaultsecurity-group>

A default security group is named default, and it has an ID assigned by AWS. The following are the initial settings for each default security group:

Allow inbound traffic only from other instances associated with the default security group Allow all outbound traffic from the instance

The default security group specifies itself as a source security group in its inbound rules. This is what allows instances associated with the default security group to communicate with other instances associated with the default security group.

**Default Security Groups**

Your AWS account automatically has a *default security group* per VPC and per region for EC2-Classical. If you don't specify a security group when you launch an instance, the instance is automatically associated with the default security group.

A default security group is named `default`, and it has an ID assigned by AWS. The following are the default rules for each default security group:

- Allows all inbound traffic from other instances associated with the default security group (the security group specifies itself as a source security group in its inbound rules)
- Allows all outbound traffic from the instance.

You can add or remove the inbound rules for any default security group. You can add or remove outbound rules for any VPC default security group.

You can't delete a default security group. If you try to delete the EC2-Classical default security group, you'll get the following error: `Client.InvalidGroup.Reserved: The security group 'default' is reserved.` If you try to delete a VPC default security group, you'll get the following error: `Client.CannotDelete: the specified group: "sg-51530134" name: "default" cannot be deleted by a user.`

**NEW QUESTION 168**

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

**Explanation:**

<https://aws.amazon.com/route53/>

#### NEW QUESTION 171

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

**Explanation:**

"If an EBS volume is the root device of an instance, you must stop the instance before you can detach the volume."

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-detaching-volume.html>

#### NEW QUESTION 176

Every user you create in the IAM system starts with \_\_\_\_\_.

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

**Answer:** B

**Explanation:**

Permissions let you specify who has access to AWS resources, and what actions they can perform on those resources.

Every IAM user starts with no permissions. [http://docs.aws.amazon.com/IAM/latest/UserGuide/access\\_permissions.html#NoDefaultPermission](http://docs.aws.amazon.com/IAM/latest/UserGuide/access_permissions.html#NoDefaultPermission)

#### NEW QUESTION 180

Which of the following statements are true about Amazon Route 53 resource records? (Choose two.)

- A. An Alias record can map one DNS name to another Amazon Route 53 DNS name.
- B. A CNAME record can be created for your zone apex.
- C. An Amazon Route 53 CNAME record can point to any DNS record hosted anywhere.
- D. TTL can be set for an Alias record in Amazon Route 53.
- E. An Amazon Route 53 Alias record can point to any DNS record hosted anywhere

**Answer:** AC

#### NEW QUESTION 181

A \_\_\_\_\_ is an individual, system, or application that interacts with AWS programmatically.

- A. user
- B. AWS Account
- C. Group
- D. Role

**Answer:** A

**Explanation:**

Q: What is a user?

A user is a unique identity recognized by AWS services and applications. Similar to a login user in an operating system like Windows or UNIX, a user has a unique name and can identify itself

using familiar security credentials such as a password or access key. A user can be an individual, system, or application requiring access to AWS services. IAM supports users (referred to as "IAM users") managed in AWS's identity management system, and it also enables you to grant access to AWS resources for users managed outside of AWS in your corporate directory (referred to as "federated users").

#### NEW QUESTION 185

Select the correct statement:

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

**Answer:** C

**Explanation:**

You can't terminate, stop, or delete a resource based solely on its tags; you must specify the resource identifier.  
[http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using\\_Tags.html#tag-restrictions](http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html#tag-restrictions)

**NEW QUESTION 190**

Which AWS instance address has the following characteristics? : "If you stop an instance, its Elastic IP address is unmapped, and you must remap it when you restart the instance."

- A. VPC Addresses
- B. EC2 Addresses
- C. Both A and B
- D. None of the above

**Answer: B**

**NEW QUESTION 192**

True or False: Common points of failures like generators and cooling equipment are shared across Availability Zones.

- A. TRUE
- B. FALSE

**Answer: B**

**NEW QUESTION 195**

Multi-AZ deployment \_\_\_\_\_ supported for Microsoft SQL Server DB Instances.

- A. is not currently
- B. is as of 2013
- C. is planned to be in 2014
- D. will never be

**Answer: C**

**NEW QUESTION 199**

Because of the extensibility limitations of striped storage attached to Windows Server, Amazon RDS does not currently support increasing storage on a \_\_\_\_\_ DB Instance.

- A. SQL Server
- B. MySQL
- C. Oracle

**Answer: A**

**NEW QUESTION 204**

Can the string value of 'Key' be prefixed with laws?

- A. No
- B. Only for EC2 not S3
- C. Yes
- D. Only for S3 not EC

**Answer: A**

**NEW QUESTION 209**

Are you able to integrate a multi-factor token service with the AWS Platform?

- A. Yes, you can integrate private multi-factor token devices to authenticate users to the AWS platform.
- B. No, you cannot integrate multi-factor token devices with the AWS platform.
- C. Yes, using the AWS multi-factor token devices to authenticate users on the AWS platform

**Answer: C**

**Explanation:**

Private MFA does not apply here.

Q. What is AWS MFA?

AWS multi-factor authentication (AWS MFA) provides an extra level of security that you can apply to your AWS environment. You can enable AWS MFA for your AWS account and for individual AWS Identity and Access Management (IAM) users you create under your account.

**NEW QUESTION 211**

Fill in the blanks: \_\_\_\_\_ is a durable, block-level storage volume that you can attach to a single, running Amazon EC2 instance.

- A. Amazon S3
- B. Amazon EBS
- C. None of these
- D. All of these

**Answer: B**

**NEW QUESTION 216**

What does Amazon RDS stand for?

- A. Regional Data Server.
- B. Relational Database Service.
- C. Nothing.
- D. Regional Database Serv

**Answer: B**

**NEW QUESTION 217**

What does Amazon ELB stand for?

- A. Elastic Linux Box.
- B. Encrypted Linux Box.
- C. Encrypted Load Balancing.
- D. Elastic Load Balancin

**Answer: D**

**NEW QUESTION 222**

Location of Instances is \_\_\_\_\_

- A. Regional
- B. based on Availability Zone
- C. Global

**Answer: B**

**Explanation:**

Regions and Availability Zones

Amazon EC2 is hosted in multiple locations world-wide. These locations are composed of regions and Availability Zones. Each region is a separate geographic area. Each region has multiple, isolated locations known as Availability Zones. Amazon EC2 provides you the ability to place resources, such as instances, and data in multiple locations. Resources aren't replicated across regions unless you do so specifically. <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availabilityzones.html#concepts-regions-availability-zones>

```
$ aws ec2 describe-availability-zones --region us-east-1
{
  "AvailabilityZones": [
    {
      "State": "available",
      "RegionName": "us-east-1",
      "Messages": [],
      "ZoneName": "us-east-1b"
    },
    {
      "State": "available",
      "RegionName": "us-east-1",
      "Messages": [],
      "ZoneName": "us-east-1c"
    },
    {
      "State": "available",
      "RegionName": "us-east-1",
      "Messages": [],
      "ZoneName": "us-east-1d"
    }
  ]
}
```

**NEW QUESTION 225**

Does Dynamic DB support in-place atomic updates?

- A. It is not defined
- B. No
- C. Yes
- D. It does support in-place non-atomic updates

**Answer: C**

**Explanation:**

Q: Does DynamoDB support in-place atomic updates?

Amazon DynamoDB supports fast in-place updates. You can increment or decrement a numeric attribute in a row using a single API call. Similarly, you can atomically add or remove to sets, lists, or maps.

<https://aws.amazon.com/dynamodb/faqs/>

**NEW QUESTION 228**

Can I attach more than one policy to a particular entity?

- A. Yes always

- B. Only if within GovCloud
- C. No
- D. Only if within VPC

**Answer:** A

**NEW QUESTION 230**

Fill in the blanks: A \_\_\_\_\_ is a storage device that moves data in sequences of bytes or bits (blocks). Hint: These devices support random access and generally use buffered I/O.

- A. block map
- B. storage block
- C. mapping device
- D. block device

**Answer:** D

**NEW QUESTION 235**

REST or Query requests are HTTP or HTTPS requests that use an HTTP verb (such as GET or POST) and a parameter named Action or Operation that specifies the API you are calling.

- A. FALSE
- B. TRUE

**Answer:** B

**Explanation:**

<http://docs.aws.amazon.com/AWSEC2/latest/APIReference/Query-Requests.html>



**NEW QUESTION 239**

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:** C

**Explanation:**

Each AWS Direct Connect location enables connectivity to all Availability Zones within the geographically nearest AWS region.

**NEW QUESTION 243**

What does the "Server Side Encryption" option on Amazon S3 provide?

- A. It provides an encrypted virtual disk in the Cloud.
- B. It doesn't exist for Amazon S3, but only for Amazon EC2.
- C. It encrypts the files that you send to Amazon S3, on the server side.
- D. It allows to upload files using an SSL endpoint, for a secure transfe

**Answer:** C

**Explanation:**

Server-side encryption is about protecting data at rest. Server-side encryption with Amazon S3- managed encryption keys (SSE-S3) employs strong multi-factor encryption.

Amazon S3 encrypts each object with a unique key. As an additional safeguard, it encrypts the key itself with a master key that it regularly rotates. Amazon S3 server-side encryption uses one of the strongest block ciphers available, 256-bit Advanced Encryption Standard (AES-256), to encrypt your data.

References:

**NEW QUESTION 245**

Will I be charged if the DB instance is idle?

- A. No
- B. Yes
- C. Only is running in GovCloud
- D. Only if running in VPC

**Answer:** B

#### NEW QUESTION 250

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 RDMS used
- B. TRUE
- C. FALSE

**Answer:** B

#### Explanation:

**DB Security Groups**

Each DB security group rule enables a specific source to access a DB instance that is associated with that DB security group. The source can be a range of addresses (e.g., 203.0.113.0/24), or an EC2 security group. When you specify an EC2 security group as the source, you allow incoming traffic from all EC2 instances that use that EC2 security group. Note that DB security group rules apply to inbound traffic only; outbound traffic is not currently permitted for DB instances.

You do not need to specify a destination port number when you create DB security group rules; the port number defined for the DB instance is used as the destination port number for all rules defined for the DB security group. DB security groups can be created using the Amazon RDS APIs or the Amazon RDS page of the AWS Management Console.

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.RDSSecurityGroups.html>

#### NEW QUESTION 251

Can I initiate a "forced failover" for my Oracle Multi-AZ DB Instance deployment?

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

**Answer:** A

#### Explanation:

<https://aws.amazon.com/public-data-sets/>

If your DB instance is a Multi-AZ deployment, you can force a failover from one availability zone to another when you select the Reboot option. When you force a failover of your DB instance, Amazon RDS automatically switches to a standby replica in another Availability Zone and updates the DNS record for the DB instance to point to the standby DB instance. As a result, you will need to clean up and re-establish any existing connections to your DB instance. Reboot with failover is beneficial when you want to simulate a failure of a DB instance for testing, or restore operations to the original AZ after a failover occurs.

Source: [http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER\\_RebootInstance.html](http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_RebootInstance.html)

#### NEW QUESTION 254

Amazon EC2 provides a repository of public data sets that can be seamlessly integrated into AWS cloud-based applications. What is the monthly charge for using the public data sets?

- A. A 1 time charge of 10\$ for all the datasets.
- B. 1\$ per dataset per month
- C. 10\$ per month for all the datasets
- D. There is no charge for using the public data sets

**Answer:** D

#### NEW QUESTION 259

In the Amazon RDS Oracle DB engine, the Database Diagnostic Pack and the Database Tuning Pack are only available with \_\_\_\_\_

- A. Oracle Standard Edition
- B. Oracle Express Edition
- C. Oracle Enterprise Edition
- D. None of these

**Answer:** C

#### Explanation:

<https://www.pythian.com/blog/a-most-simple-cloud-is-amazon-rds-for-oracle-right-for-you/>

#### NEW QUESTION 264

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 265**

Amazon RDS supports SOAP only through \_\_\_\_\_.

- A. HTTP or HTTPS
- B. TCP/IP
- C. HTTP
- D. HTTPS

**Answer:** D

**Explanation:**

Amazon RDS supports SOAP only through HTTPS

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/using-soap-api.html>

**WSDL and Schema Definitions**

You can access the Amazon Relational Database Service using the SOAP web services messaging protocol. This interface is described by a Web Services Description Language (WSDL) document, which defines the operations and security model for the particular service. The WSDL references an XML Schema document, which strictly defines the data types that might appear in SOAP requests and responses. For more information on WSDL and SOAP, see [Web Services References](#).

**Note**

Amazon RDS supports SOAP only through HTTPS.

**NEW QUESTION 269**

The SQL Server \_\_\_\_\_ feature is an efficient means of copying data from a source database to your DB Instance. It writes the data that you specify to a data file, such as an ASCII file.

- A. bulk copy
- B. group copy
- C. dual copy
- D. mass copy

**Answer:** A

**Explanation:**

The SQL Server bulk copy feature is an efficient means of copying data from a source database to your DB Instance. Bulk copy writes the data that you specify to a data file, such as an ASCII

file. You can then run bulk copy again to write the contents of the file to the destination DB Instance.

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/SQLServer.Procedural.Importing.html>

**NEW QUESTION 271**

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

**Explanation:**

<https://aws.amazon.com/ebs/pricing/>

**NEW QUESTION 275**

You can modify the backup retention period; valid values are 0 (for no backup retention) to a maximum of \_\_\_\_\_ days.

- A. 45
- B. 35
- C. 15
- D. 5

**Answer:** B

**Explanation:**

[http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER\\_WorkingWithAutomatedBackups.html](http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithAutomatedBackups.html)

**NEW QUESTION 277**

Which of the following features ensures even distribution of traffic to Amazon EC2 instances in multiple Availability Zones registered with a load balancer?

- A. Elastic Load Balancing request routing
- B. An Amazon Route 53 weighted routing policy
- C. Elastic Load Balancing cross-zone load balancing
- D. An Amazon Route 53 latency routing policy

**Answer: C**

**Explanation:**

Cross-zone load balancing is always enabled for an Application Load Balancer and is disabled by default for a Classic Load Balancer. If cross-zone load balancing is enabled, the load balancer distributes traffic evenly across all registered instances in all enabled Availability Zones. If cross-zone load balancing is disabled, the load balancer distributes traffic evenly across all enabled Availability Zones. For example, suppose that you have 10 instances in Availability Zone us-west-2a and 2 instances in us-west-2b. If cross-zone load balancing is disabled, the requests are distributed evenly between us-west-2a and us-west-2b. As a result, the 2 instances in us-west-2b serve the same amount of traffic as the 10 instances in us-west-2a. However, if cross-zone load balancing is enabled, the load balancer distributes incoming requests evenly across all 12 instances. <http://docs.aws.amazon.com/elasticloadbalancing/latest/userguide/how-elastic-load-balancingworks.html>

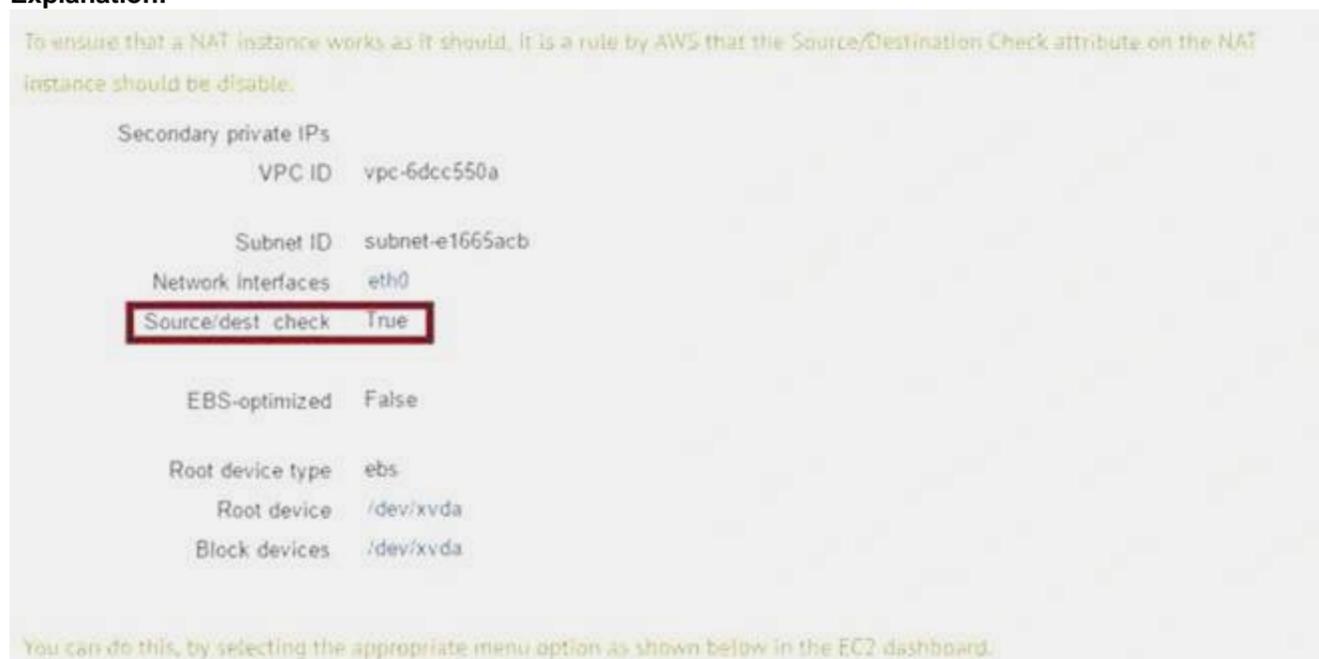
**NEW QUESTION 279**

After launching an instance that you intend to serve as a NAT (Network Address Translation) device in a public subnet you modify your route tables to have the NAT device be the target of internet bound traffic of your private subnet. When you try and make an outbound connection to the internet from an instance in the private subnet, you are not successful. Which of the following steps could resolve the issue?

- A. Disabling the Source/Destination Check attribute on the NAT instance
- B. Attaching an Elastic IP address to the instance in the private subnet
- C. Attaching a second Elastic Network Interface (ENI) to the NAT instance, and placing it in the private subnet
- D. Attaching a second Elastic Network Interface (ENI) to the instance in the private subnet, and placing it in the public subnet

**Answer: A**

**Explanation:**



**NEW QUESTION 283**

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 changes.

**Answer: A**

**Explanation:**

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html#vpc-securitygroups>

**NEW QUESTION 287**

A customer is leveraging Amazon Simple Storage Service in eu-west-1 to store static content for a web-based property. The customer is storing objects using the Standard Storage class. Where are the customer's objects replicated?

- A. A single facility in eu-west-1 and a single facility in eu-central-1
- B. A single facility in eu-west-1 and a single facility in us-east-1
- C. Multiple facilities in eu-west-1
- D. A single facility in eu-west-1

**Answer: C**

**Explanation:**

Objects stored in a region never leave the region unless you explicitly transfer them to another region. For example, objects stored in the EU (Ireland) region never leave it. <http://docs.aws.amazon.com/AmazonS3/latest/dev/Introduction.html#Regions>

**NEW QUESTION 292**

You have decided to change the instance type for instances running in your application tier that is using Auto Scaling. In which area below would you change the instance type definition?

- A. Auto Scaling policy
- B. Auto Scaling group
- C. Auto Scaling tags
- D. Auto Scaling launch configuration

**Answer: D**

**NEW QUESTION 297**

A company wants to implement their website in a virtual private cloud (VPC). The web tier will use an Auto Scaling group across multiple Availability Zones (AZs). The database will use Multi-AZ RDS MySQL and should not be publicly accessible. What is the minimum number of subnets that need to be configured in the VPC?

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

**Answer: D**

**Explanation:**

Since multi-AZ RDS needs 2 private subnets to provide high availability and 2 public subnets are needed for ELB(web-tier) application.

Would use VPC with private (DB) and public (WEB) subnets: [http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER\\_VPC.Scenarios.html](http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.Scenarios.html) Multi AZ requirement forces me to multiply subnets by two.

Reasons:

For DB: Your VPC must have at least one subnet in at least two of the Availability Zones in the region where you want to deploy your DB instance. A subnet is a segment of a VPC's IP address

range that you can specify and that lets you group instances based on your security and operational needs

[http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER\\_VPC.WorkingWithRDSInstancein aVPC.html](http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.WorkingWithRDSInstancein aVPC.html)

For Web: After creating a VPC, you can add one or more subnets in each Availability Zone. Each subnet must reside entirely within one Availability Zone and cannot span zones [http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC\\_Subnets.html](http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html)

**NEW QUESTION 301**

For which of the following use cases are Simple Workflow Service (SWF) and Amazon EC2 an appropriate solution? (Choose two.)

- A. Using as an endpoint to collect thousands of data points per hour from a distributed fleet of sensors
- B. Managing a multi-step and multi-decision checkout process of an e-commerce website
- C. Orchestrating the execution of distributed and auditable business processes
- D. Using as an SNS (Simple Notification Service) endpoint to trigger execution of video transcoding jobs
- E. Using as a distributed session store for your web application

**Answer: BC**

**Explanation:**

<https://aws.amazon.com/swf/faqs/>

**NEW QUESTION 306**

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 balancer

**Answer: B**

**Explanation:**

Elastic Load Balancing access logs

The access logs for Elastic Load Balancing capture detailed information for all requests made to your load balancer and stores them as log files in the Amazon S3 bucket that you specify. Each log contains

details such as the time a request was received, the client's IP address, latencies, request path, and server responses. You can use these access logs to analyze traffic patterns and to troubleshoot your back-end applications. For more information, see [Monitor Your Load Balancer Using Elastic Load Balancing Access Logs](#).

**NEW QUESTION 307**

A company is storing data on Amazon Simple Storage Service (S3). The company's security policy mandates that data is encrypted at rest. Which of the following methods can achieve this? (Choose three.)

- A. Use Amazon S3 server-side encryption with AWS Key Management Service managed keys.
- B. Use Amazon S3 server-side encryption with customer-provided keys.
- C. Use Amazon S3 server-side encryption with EC2 key pair.

- D. Use Amazon S3 bucket policies to restrict access to the data at rest.
- E. Encrypt the data on the client-side before ingesting to Amazon S3 using their own master key.
- F. Use SSL to encrypt the data while in transit to Amazon S3.

**Answer:** ABE

**NEW QUESTION 312**

Which of the following are valid statements about Amazon S3? (Choose two.)

- A. S3 provides read-after-write consistency for any type of PUT or DELETE
- B. Consistency is not guaranteed for any type of PUT or DELETE
- C. A successful response to a PUT request only occurs when a complete object is saved.
- D. Partially saved objects are immediately readable with a GET after an overwrite PUT.
- E. S3 provides eventual consistency for overwrite PUTS and DELETE

**Answer:** CE

**Explanation:**



**NEW QUESTION 316**

Which features can be used to restrict access to data in S3? (Choose two.)

- A. Set an S3 ACL on the bucket or the object.
- B. Create a CloudFront distribution for the bucket.
- C. Set an S3 bucket policy.
- D. Enable IAM Identity Federation
- E. Use S3 Virtual Hosting

**Answer:** AC

**Explanation:**

Amazon S3 is secure by default. Only the bucket and object owners originally have access to Amazon S3 resources they create. Amazon S3 supports user authentication to control access to data. You can use access control mechanisms such as bucket policies and Access Control Lists (ACLs) to selectively grant permissions to users and groups of users. You can securely upload/download your data to Amazon S3 via SSL endpoints using the HTTPS protocol. If you need extra security you can use the Server Side Encryption (SSE) option or the Server Side Encryption with Customer-Provide Keys (SSEC) option to encrypt data stored-at-rest. Amazon S3 provides the encryption technology for both SSE and SSE-C. Alternatively you can use your own encryption libraries to encrypt data before storing it in Amazon S3.  
<https://aws.amazon.com/s3/faqs/>

**NEW QUESTION 317**

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

**NEW QUESTION 321**

You have a distributed application that periodically processes large volumes of data across multiple Amazon EC2 Instances. The application is designed to recover gracefully from Amazon EC2 instance failures. You are required to accomplish this task in the most cost-effective way. Which of the following will meet your requirements?

- A. Spot Instances
- B. Reserved instances
- C. Dedicated instances
- D. On-Demand instances

**Answer:** A

**Explanation:**

Using reserved instances is not the most cost-effective way. <https://aws.amazon.com/blogs/aws/new-scheduled-reserved-instances/>  
"Scheduled Reserved

Instance model allows you to reserve instances for predefined blocks of time on a recurring basis for a one-year term, with prices that are generally 5 to 10% lower than the equivalent On-Demand rates." You can get spot instances with much lower prices: <https://aws.amazon.com/ec2/spot/pricing/>

"Spot instances are also available to run for a predefined duration in hourly increments up to six hours in length at a significant discount (30-45%) compared to On-Demand pricing plus an additional 5% during off-peak times for a total of up to 50% savings."

**NEW QUESTION 326**

You have a content management system running on an Amazon EC2 instance that is approaching 100% CPU utilization. Which option will reduce load on the Amazon EC2 instance?

- A. Create a load balancer, and register the Amazon EC2 instance with it
- B. Create a CloudFront distribution, and configure the Amazon EC2 instance as the origin
- C. Create an Auto Scaling group from the instance using the CreateAutoScalingGroup action
- D. Create a launch configuration from the instance using the CreateLaunchConfiguration action

**Answer: C**

**Explanation:**

You can create an ASG from instance ID [http://docs.aws.amazon.com/AutoScaling/latest/APIReference/API\\_CreateAutoScalingGroup.html](http://docs.aws.amazon.com/AutoScaling/latest/APIReference/API_CreateAutoScalingGroup.html)

**NEW QUESTION 331**

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 balancer's DNS name. Which options are probable causes of this behavior? (Choose two.)

- A. The load balancer was not configured to use a public subnet 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 properly 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**

**Explanation:**

There is no such thing as VGW. Hence E is not the correct answer.

**NEW QUESTION 334**

A company needs to deploy services to an AWS region which they have not previously used. The company currently has an AWS identity and Access Management (IAM) role for the Amazon EC2 instances, which permits the instance to have access to Amazon DynamoDB. The company wants their EC2 instances in the new region to have the same privileges. How should the company achieve this?

- A. Create a new IAM role and associated policies within the new region
- B. Assign the existing IAM role to the Amazon EC2 instances in the new region
- C. Copy the IAM role and associated policies to the new region and attach it to the instances
- D. Create an Amazon Machine Image (AMI) of the instance and copy it to the desired region using the AMI Copy feature

**Answer: B**

**NEW QUESTION 339**

Which of the following notification endpoints or clients are supported by Amazon Simple Notification Service? (Choose two.)

- A. Email
- B. CloudFront distribution
- C. File Transfer Protocol
- D. Short Message Service
- E. Simple Network Management Protocol

**Answer: AD**

**NEW QUESTION 341**

You are deploying an application to track GPS coordinates of delivery trucks in the United States. Coordinates are transmitted from each delivery truck 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 Service

**Answer: A**

**Explanation:**

<https://aws.amazon.com/streaming-data/>

**NEW QUESTION 345**

You have an application running on an Amazon Elastic Compute Cloud instance, that uploads 5 GB video objects to Amazon Simple Storage Service (S3). 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 S3 multipart upload
- C. Leveraging Amazon CloudFront, use the HTTP POST method to reduce latency.
- D. Use Amazon Elastic Block Store Provisioned IOPs and use an Amazon EBS-optimized instance

**Answer: B**

**Explanation:**

Using multipart upload provides the following advantages:

- Improved throughput - You can upload parts in parallel to improve throughput.
- Quick recovery from any network issues - Smaller part size minimizes the impact of restarting a failed upload due to a network error.
- Pause and resume object uploads - You can upload object parts over time. Once you initiate a multipart upload there is no expiry; you must explicitly complete or abort the multipart upload.
- Begin an upload before you know the final object size.
- You can upload an object as you are creating it. <http://docs.aws.amazon.com/AmazonS3/latest/dev/uploadobjusingmpu.html>

**NEW QUESTION 349**

A customer wants to track access to their Amazon Simple Storage Service (S3) buckets and also use this information for their internal security and access audits. Which of the following will meet the Customer requirement?

- A. Enable AWS CloudTrail to audit all Amazon S3 bucket access.
- B. Enable server access logging for all required Amazon S3 buckets.
- C. Enable the Requester Pays option to track access via AWS Billing
- D. Enable Amazon S3 event notifications for Put and Pos

**Answer: B**

**Explanation:**

References:

<http://docs.aws.amazon.com/AmazonS3/latest/dev/ServerLogs.html> <http://docs.aws.amazon.com/AmazonS3/latest/dev/cloudtrail-logging.html>

**NEW QUESTION 351**

A customer is hosting their company website on a cluster of web servers that are behind a publicfacing load balancer. The customer also uses Amazon Route 53 to manage their public DNS. How should the customer configure the DNS zone apex record to point to the load balancer?

- A. Create an A record pointing to the IP address of the load balancer
- B. Create a CNAME record pointing to the load balancer DNS name.
- C. Create a CNAME record aliased to the load balancer DNS name.
- D. Create an A record aliased to the load balancer DNS name

**Answer: D**

**NEW QUESTION 352**

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 establish

**Answer: AC**

**NEW QUESTION 357**

You are building an automated transcription service in which Amazon EC2 worker instances process an uploaded audio file and generate a text file. You must store both of these files in the same durable storage until the text file is retrieved. You do not know what the storage capacity requirements are. Which storage option is both cost-efficient and scalable?

- A. Multiple Amazon EBS volume with snapshots
- B. A single Amazon Glacier vault
- C. A single Amazon S3 bucket
- D. Multiple instance stores

**Answer: C**

**NEW QUESTION 360**

A company is deploying a new two-tier web application in AWS. The company has limited staff and requires high availability, and the application requires complex queries and table joins. Which configuration provides the solution for the company's requirements?

- A. MySQL Installed on two Amazon EC2 Instances in a single Availability Zone
- B. Amazon RDS for MySQL with Multi-AZ

- C. Amazon ElastiCache
- D. Amazon DynamoDB

**Answer:** B

**Explanation:**

When is it appropriate to use DynamoDB instead of a relational database? From our own experience designing and operating a highly available, highly scalable ecommerce platform, we have come to realize that relational databases should only be used when an application really needs the complex query, table join and transaction capabilities of a full-blown relational database. In all other cases, when such relational features are not needed, a NoSQL database service like DynamoDB offers a simpler, more available, more scalable and ultimately a lower cost solution.

**NEW QUESTION 361**

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:** C

**Explanation:**

You must launch a T2 instance using an HVM AMI. For more information, see Linux AMI Virtualization Types. You must launch your T2 instances using an EBS volume as the root device. For more information, see Amazon EC2 Root Device Volume. <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-resize.html>

**NEW QUESTION 362**

You manually launch a NAT AMI in a public subnet. The network is properly configured. Security groups and network access control lists are properly configured. Instances in a private subnet can access the NAT. The NAT can access the Internet. However, private instances cannot access the Internet. What additional step is required to allow access from the private instances?

- A. Enable Source/Destination Check on the private Instances.
- B. Enable Source/Destination Check on the NAT instance.
- C. Disable Source/Destination Check on the private instances.
- D. Disable Source/Destination Check on the NAT instance

**Answer:** D

**Explanation:**

Disabling Source/Destination Checks.

Each EC2 instance performs source/destination checks by default. This means that the instance must be the source or destination of any traffic it sends or receives. However, a NAT instance must be able to send and receive traffic when the source or destination is not itself. Therefore, you must disable source/destination checks on the NAT instance. You can disable the SrcDestCheck attribute for a NAT instance that's either running or stopped using the console or the command line. [http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC\\_NAT\\_Instance.html](http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_NAT_Instance.html)

**NEW QUESTION 364**

Which of the following approaches provides the lowest cost for Amazon Elastic Block Store snapshots while giving you the ability to fully restore data?

- A. Maintain two snapshots: the original snapshot and the latest incremental snapshot.
- B. Maintain a volume snapshot; subsequent snapshots will overwrite one another
- C. Maintain a single snapshot the latest snapshot is both Incremental and complete.
- D. Maintain the most current snapshot, archive the original and incremental to Amazon Glacie

**Answer:** C

**Explanation:**

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-deleting-snapshot.html>

**NEW QUESTION 368**

A company has configured and peered two VPCs: VPC-1 and VPC-2. VPC-1 contains only private subnets, and VPC-2 contains only public subnets. The company uses a single AWS Direct Connect connection and private virtual interface to connect their on-premises network with VPC-1. Which two methods increases the fault tolerance of the connection to VPC-1? (Choose two.)

- A. Establish a hardware VPN over the internet between VPC-2 and the on-premises network.
- B. Establish a hardware VPN over the internet between VPC-1 and the on-premises network.
- C. Establish a new AWS Direct Connect connection and private virtual interface in the same region as VPC-2.
- D. Establish a new AWS Direct Connect connection and private virtual interface in a different AWS region than VPC-1.
- E. Establish a new AWS Direct Connect connection and private virtual interface in the same AWS region as VPC-1

**Answer:** BE

**NEW QUESTION 372**

You work for a major news network in Europe. They have just released a new app which allows users to report on events as and when they happen using their mobile phone. Users are able to upload pictures from the app and then other users will be able to view these pics. Your organization expects this app to grow very quickly, essentially doubling it's user base every month. The app uses S3 to store the media and you are expecting sudden and large increases in traffic to S3 when a major news event takes place (as people will be uploading content in huge numbers). You need to keep your storage costs to a minimum however and it does not matter if some objects are lost. Which storage media should you use to keep costs as low as possible?

- A. S3 - Infrequently Accessed Storage.
- B. S3 - Reduced Redundancy Storage (RRS).
- C. Glacier.
- D. S3 - Provisioned IOP

**Answer: B**

#### NEW QUESTION 373

Your company has decided to set up a new AWS account for test and dev purposes. They already use AWS for production, but would like a new account dedicated for test and dev so as to not accidentally break the production environment. You launch an exact replica of your production environment using a CloudFormation template that your company uses in production. However CloudFormation fails. You use the exact same CloudFormation template in production, so the failure is something to do with your new AWS account. The CloudFormation template is trying to launch 60 new EC2 instances in a single AZ. After some research you discover that the problem is;

- A. For all new AWS accounts there is a soft limit of 20 EC2 instances per regio
- B. You should submit the limit increase form and retry the template after your limit has been increased.
- C. For all new AWS accounts there is a soft limit of 20 EC2 instances per availability zon
- D. You should submit the limit increase form and retry the template after your limit has been increased.
- E. You cannot launch more than 20 instances in your default VPC, instead reconfigure the CloudFormation template to provision the instances in a custom VPC.
- F. Your CloudFormation template is configured to use the parent account and not the new accoun
- G. Change the account number in the CloudFormation template and relaunch the template.

**Answer: A**

#### NEW QUESTION 376

You are a solutions architect working for a biotech company who is pioneering research in immunotherapy. They have developed a new cancer treatment that may be able to cure up to 94% of cancers. They store their research data on S3, however recently an intern accidentally deleted some critical files. You've been asked to prevent this from happening in the future. What options below can prevent this?

- A. Make sure the interns can only access data on S3 using signed URLs.
- B. Enable S3 versioning on the bucket & enable Enable Multifactor Authentication (MFA) on the bucket.
- C. Use S3 Infrequently Accessed storage to store the data on.
- D. Create an IAM bucket policy that disables delete

**Answer: B**

#### NEW QUESTION 381

You have been asked to identify a service on AWS that is a durable key value store. Which of the services below meets this definition?

- A. Mobile Hub
- B. Kinesis
- C. Simple Storage Service (S3)
- D. Elastic File Service (EFS)

**Answer: C**

#### NEW QUESTION 382

You work for a cosmetic company which has their production website on AWS. The site itself is in a two-tier configuration with web servers in the front end and database servers at the back end. The site uses using Elastic Load Balancing and Auto Scaling. The databases maintain consistency by replicating changes to each other as and when they occur. This requires the databases to have extremely low latency. Your website needs to be highly redundant and must be designed so that if one availability zone goes offline and Auto Scaling cannot launch new instances in the remaining Availability Zones the site will not go offline. How can the current architecture be enhanced to ensure this?

- A. Deploy your site in three different AZ's within the same regio
- B. Configure the Auto Scaling minimum to handle 50 percent of the peak load per zone.
- C. Deploy your website in 2 different region
- D. Configure Route53 with a failover routing policy and set up health checks on the primary site.
- E. Deploy your site in three different AZ's within the same regio
- F. Configure the Auto Scaling minimum to handle 33 percent of the peak load per zone.
- G. Deploy your website in 2 different region
- H. Configure Route53 with Weighted Routin
- I. Assign a weight of 25% to region 1 and a weight of 75% to region 2.

**Answer: A**

#### NEW QUESTION 384

You have been asked to create VPC for your company. The VPC must support both Internet-facing web applications (ie they need to be publicly accessible) and internal private applications (i.e. they are not publicly accessible and can be accessed only over VPN). The internal private applications must be inside a private subnet. Both the internet-facing and private applications must be able to leverage at least three Availability Zones for high availability. At a minimum, how many subnets must you create within your VPC to achieve this?

- A. 5
- B. 3
- C. 4
- D. 6

**Answer: D**

**NEW QUESTION 385**

Amazon's Redshift uses which block size for its columnar storage?

- A. 2KB
- B. 8KB
- C. 16KB
- D. 32KB
- E. 1024KB / 1MB

**Answer:** E

**NEW QUESTION 387**

When creating an RDS instance you can select which availability zone in which to deploy your instance.

- A. True
- B. False

**Answer:** A

**NEW QUESTION 392**

Amazon S3 buckets in all other regions (other than US Standard) do not provide eventual consistency for overwrite PUTS and DELETES.

- A. True
- B. False

**Answer:** B

**NEW QUESTION 396**

It is possible to transfer a reserved instance from one Availability Zone to another.

- A. True
- B. False

**Answer:** A

**NEW QUESTION 399**

Amazon S3 buckets in the US Standard region do not provide eventual consistency.

- A. True
- B. False

**Answer:** B

**NEW QUESTION 400**

What are the different types of virtualization available on EC2?

- A. Pseudo-Virtual (PV) & Hardware Virtual Module (HSM)
- B. Para-Virtual (PV) & Hardware Virtual Machine (HVM)
- C. Pseudo-Virtual (PV) & Hardware Virtual Machine (HVM)
- D. Para-Virtual (PV) & Hardware Virtual Module (HSM)

**Answer:** B

**NEW QUESTION 404**

Which of the following is not supported by AWS Import/Export?

- A. Import to Amazon S3
- B. Export from Amazon S3
- C. Import to Amazon EBS
- D. Import to Amazon Glacier
- E. Export to Amazon Glacier

**Answer:** E

**Explanation:**

**Overview of AWS Import/Export**

AWS Import/Export accelerates transferring data between the AWS cloud and portable storage devices that you mail to us. AWS Import/Export is a good choice if you have 16 terabytes (TB) or less of data to import into Amazon Simple Storage Service (Amazon S3), Amazon Glacier, or Amazon Elastic Block Store (Amazon EBS). You can also export data from Amazon S3 with AWS Import/Export.

#### NEW QUESTION 409

What are the valid methodologies for encrypting data on S3?

- A. Server Side Encryption (SSE)-S3, SSE-C, SSE-KMS or a client library such as Amazon S3 Encryption Client.
- B. Server Side Encryption (SSE)-S3, SSE-A, SSE-KMS or a client library such as Amazon S3 Encryption Client.
- C. Server Side Encryption (SSE)-S3, SSE-C, SSE-SSL or a client library such as Amazon S3 Encryption Client.
- D. Server Side Encryption (SSE)-S3, SSE-C, SSE-SSL or a server library such as Amazon S3 Encryption Client.

**Answer:** A

#### NEW QUESTION 410

A Provisioned IOPS SSD volume must be at least \_\_\_\_ GB in size.

- A. 1
- B. 6
- C. 20
- D. 4

**Answer:** D

#### NEW QUESTION 414

Is there a method or command in the IAM system to allow or deny access to a specific instance?

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

**Answer:** B

#### Explanation:

[http://docs.aws.amazon.com/IAM/latest/UserGuide/reference\\_policies\\_evaluationlogic.html#policy-eval-denyallow](http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_evaluationlogic.html#policy-eval-denyallow)

- By default, all requests are denied. (In general, requests made using the account credentials for resources in the account are always allowed.)
- An explicit allow overrides this default.
- An explicit deny overrides any allows.

#### NEW QUESTION 417

What is the default VPC security group limit?

- A. 500
- B. 50
- C. 5
- D. There is no limit

**Answer:** A

#### NEW QUESTION 419

How many relational database engines does RDS currently support?

- A. Three: MySQL, Oracle and Microsoft SQL Server.
- B. Just two: MySQL and Oracle.
- C. Six: Amazon Aurora, Oracle, Microsoft SQL Server, PostgreSQL, MySQL and MariaDB.
- D. Just one: MySQL

**Answer:** C

#### Explanation:

Amazon RDS provides you six familiar database engines to choose from, including Amazon Aurora, Oracle, Microsoft SQL Server, PostgreSQL, MySQL and MariaDB. <https://aws.amazon.com/rds/?nc1=hls>

Amazon Relational Database Service (Amazon RDS) makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity while managing time-consuming database administration tasks, freeing you up to focus on your applications and business. Amazon RDS provides you six familiar database engines to choose from, including Amazon Aurora, PostgreSQL, MySQL, MariaDB, Oracle, and Microsoft SQL Server.

#### NEW QUESTION 424

New database versions will automatically be applied to AWS RDS instances as they become available.

- A. True
- B. False

**Answer:** B

#### NEW QUESTION 427

Out of the striping options available for the EBS volumes, which one has the following disadvantage : 'Doubles the amount of I/O required from the instance to EBS compared to RAID 0, because you're mirroring all writes to a pair of volumes, limiting how much you can stripe.' ?

- A. Raid 5
- B. Raid 6
- C. Raid 1
- D. Raid 2

**Answer:** C

#### NEW QUESTION 431

Can an EBS volume be attached to more than one EC2 instance at the same time?

- A. No
- B. Yes.
- C. Only EC2-optimized EBS volumes.
- D. Only in read mod

**Answer:** A

#### Explanation:

EBS is network attached storage that can only be attached to one instance at a time <https://aws.amazon.com/ebs/getting-started/>  
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumes.html>

#### NEW QUESTION 432

In the basic monitoring package for EC2, Amazon CloudWatch provides the following metrics:

- A. web server visible metrics such as number failed transaction requests
- B. operating system visible metrics such as memory utilization
- C. database visible metrics such as number of connections
- D. hypervisor visible metrics such as CPU utilization

**Answer:** D

#### NEW QUESTION 436

You are building a system to distribute confidential training videos to employees. Using CloudFront, what method could be used to serve content that is stored in S3, but not publicly accessible from S3 directly?

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

**Answer:** A

#### NEW QUESTION 439

Which is an operational process performed by AWS for data security?

- A. AES-256 encryption of data stored on any shared storage device
- B. Decommissioning of storage devices using industry-standard practices
- C. Background virus scans of EBS volumes and EBS snapshots
- D. Replication of data across multiple AWS Regions
- E. Secure wiping of EBS data when an EBS volume is unmounted

**Answer:** B

#### Explanation:

"When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals.  
AWS uses the techniques detailed in DoD 5220.22-M (National Industrial Security Program Operating Manual) or NIST 800-88 (Guidelines for Media Sanitization) to destroy data as part of the decommissioning process.  
All decommissioned magnetic storage devices are degaussed and physically destroyed in accordance with industry-standard practices."

#### NEW QUESTION 441

A startup company hired you to help them build a mobile application, that will ultimately store billions of images and videos in S3. The company is lean on funding, and wants to minimize operational costs, however, they have an aggressive marketing plan, and expect to double their current installation base every six months. Due to the nature of their business, they are expecting sudden and large increases in traffic to and from S3, and need to ensure that it can handle the performance needs of their application. What other information must you gather from this customer in order to determine whether S3 is the right option?

- A. You must know how many customers the company has today, because this is critical in understanding what their customer base will be in two years.
- B. You must find out the total number of requests per second at peak usage.
- C. You must know the size of the individual objects being written to S3, in order to properly design the key namespace.
- D. In order to build the key namespace correctly, you must understand the total amount of storage needs for each S3 bucket.

**Answer:** B

#### NEW QUESTION 445

When it comes to API credentials, what is the best practice recommended by AWS?

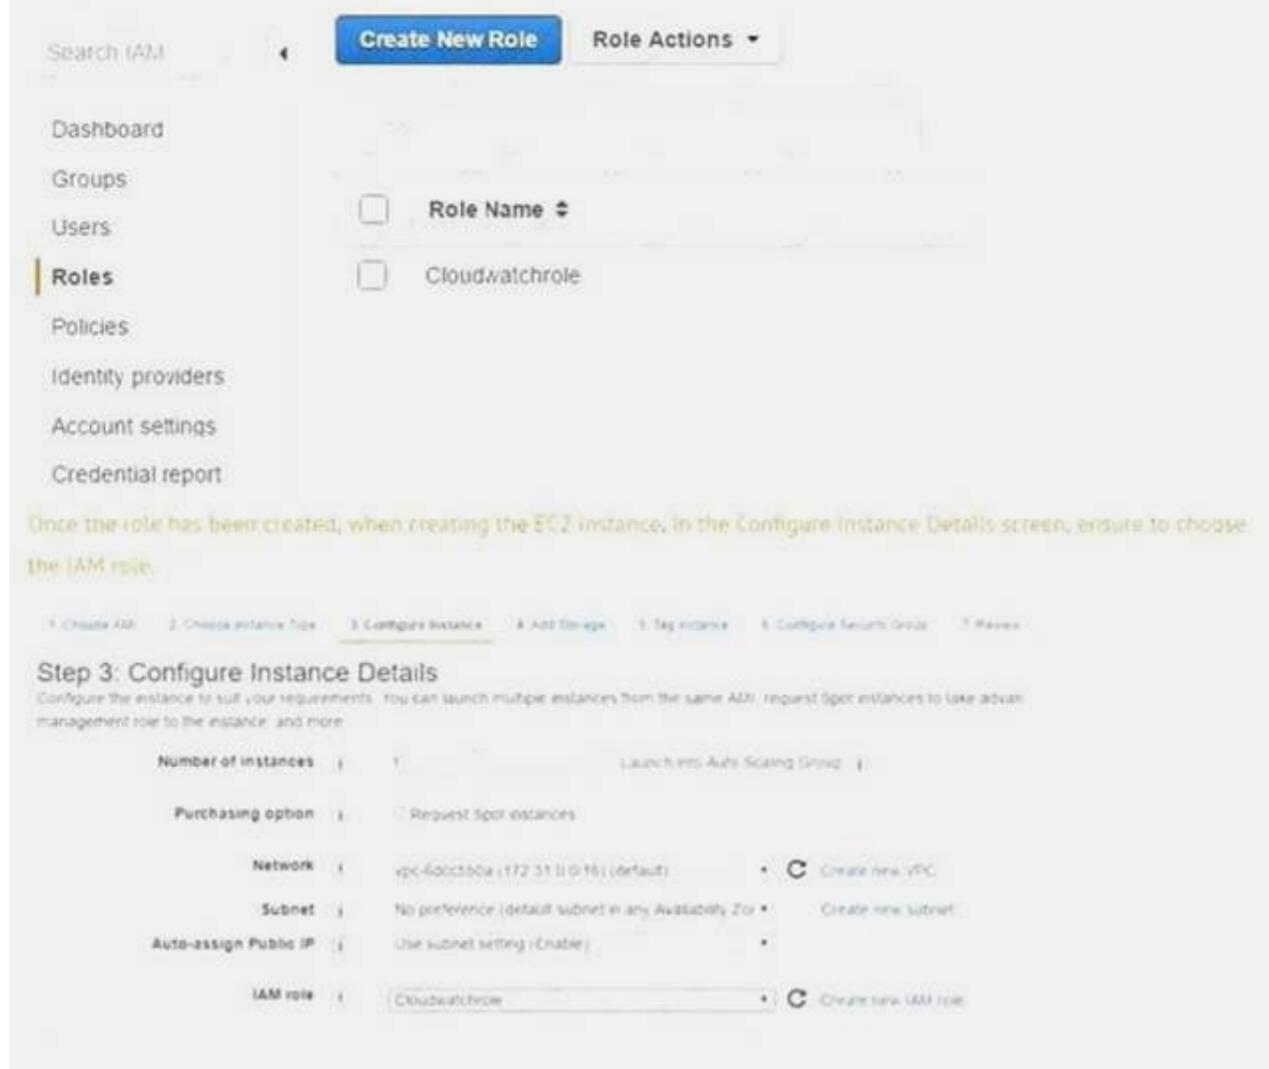
- A. Create a role which has the necessary and can be assumed by the EC2 instance.
- B. Use the API credentials from an EC2 instance.
- C. Use the API credentials from a bastion host.
- D. Use the API credentials from a NAT Instance

**Answer: A**

**Explanation:**

The best practise highlighted by AWS is always create a role which has select permissions and when creating an EC2 instance, ensure the role is attached to the EC2 instance.

So in the Security credentials in AWS, you first need to go to the Security Credentials section and create a role. The below example shows the creation of a Cloudwatch role which has the permissions to publish to cloudwatch.



Once the role has been created, when creating the EC2 instance, in the Configure Instance Details screen, ensure to choose the IAM role.

Step 3: Configure Instance Details  
 Configure the instance to suit your requirements. You can launch multiple instances from the same AMI, request Spot instances to take advantage of lower prices, request a dedicated instance to get more consistent performance, and more.

Number of instances: 1 (Launch into Auto Scaling Group)

Purchasing option:  Request Spot instances

Network: vpc-400c350a (172.31.0.0/16) (default) • Create new VPC

Subnet: No preference (default subnet in any Availability Zone) • Create new subnet

Auto-assign Public IP: Use subnet setting (Enable) •

IAM role: Cloudwatchrole • Create new IAM role

The correct answer is: Create a role which has the necessary privileges and can be assumed by the EC2 instance.

**NEW QUESTION 448**

A customer has a requirement to extend their on-premises data center to AWS. The customer requires a 50-Mbps dedicated and private connection to their VPC. Which AWS product or feature satisfies this requirement?

- A. Amazon VPC Peering
- B. Elastic IP Addresses
- C. AWS Direct Connect
- D. Amazon VPC virtual private gateway

**Answer: C**

**Explanation:**

AWS Direct connect is the solution officially provided by AWS when the customer wants to have a dedicated and private connection to their AWS cloud.

The correct answer is: AWS Direct Connect

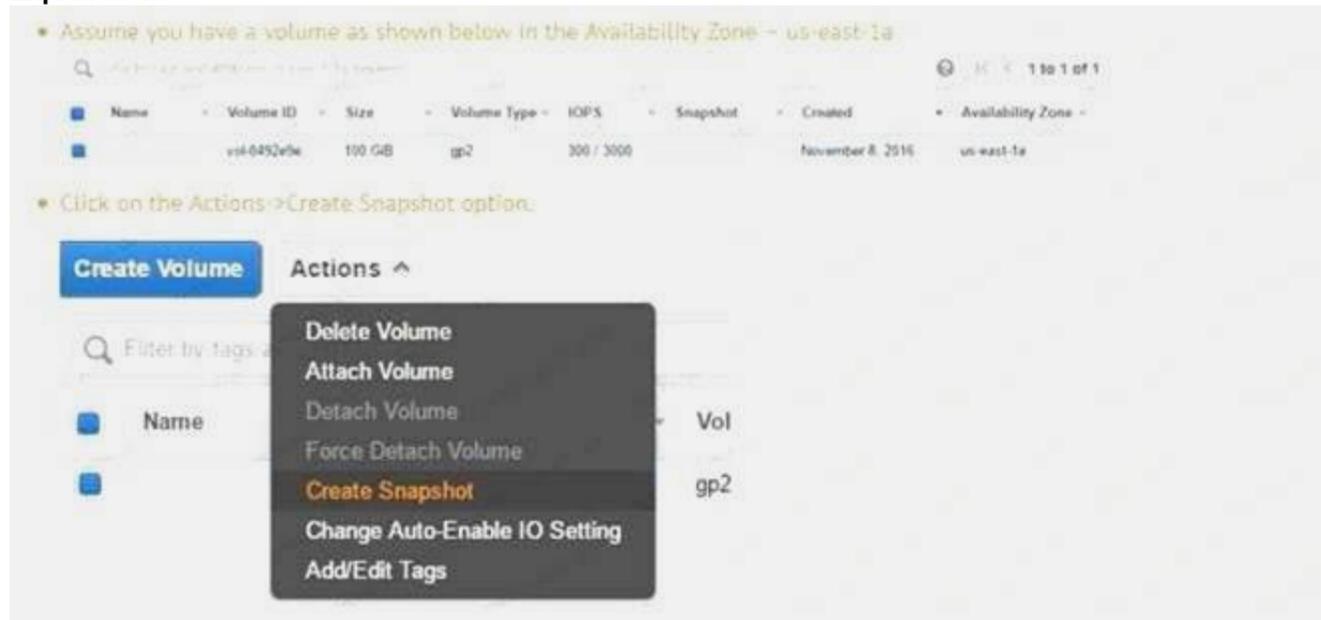
**NEW QUESTION 449**

How can an EBS volume which is currently attached to an EC2 instance in one Availability Zone to another?

- A. Detach the volume and attach to an EC2 instance in another AZ.
- B. Create a new volume in the other AZ and specify the current volume as the source.
- C. Create a snapshot of the volume and then create a volume from the snapshot in the other AZ
- D. Create a new volume in the AZ and do a disk copy of contents from one volume to another

**Answer: C**

**Explanation:**



"Snapshots can be used to instantiate multiple new volumes, expand the size of a volume, or move volumes across Availability Zones. When a new volume is created, you may choose to create it based on an existing Amazon EBS snapshot. In that scenario, the new volume begins as an exact replica of the snapshot."  
<https://aws.amazon.com/ebs/details/>

**NEW QUESTION 451**

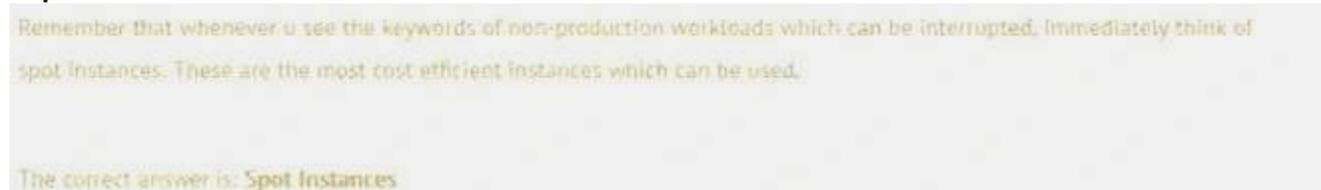
A company is hosting EC2 instances which focuses on work-loads are on non-production and nonpriority batch loads. Also these processes can be interrupted at any time.

What is the best pricing model which can be used for EC2 instances in this case?

- A. Reserved Instances
- B. On-Demand Instances
- C. Spot Instances
- D. Regular Instances

**Answer: C**

**Explanation:**



**NEW QUESTION 453**

You need a solution to distribute traffic evenly across all of the containers for a task running on Amazon ECS. Your task definitions define dynamic host port mapping for your containers. What AWS feature provides this functionality?

- A. Application Load Balancers support dynamic host port mapping.
- B. CloudFront custom origins support dynamic host port mapping.
- C. All Elastic Load Balancing instances support dynamic host port mapping.
- D. Classic Load Balancers support dynamic host port mapping.

**Answer: A**

**NEW QUESTION 455**

Your Amazon EC2 instances must access the AWS API, so you created a NAT gateway in an existing subnet. When you try to access the AWS API, you are unsuccessful.

What could be preventing access?

- A. The NAT gateway subnet does not have a route to an Internet gateway.
- B. The instances need an IAM granting access to the NAT gateway.
- C. The NAT gateway does not have a route to the virtual private gateway.
- D. The instances are not in the same subnet as the NAT gateway.

**Answer: A**

**NEW QUESTION 460**

You've been tasked with choosing a datastore to persist GPS coordinates for a new app. The service needs consistent, single-digit-millisecond latency at any scale. Which AWS service meets your requirements?

- A. Amazon Redshift
- B. Amazon DynamoDB
- C. Amazon S3
- D. Amazon RDS

**Answer: B**

**NEW QUESTION 462**

You're building an API backend available at `services.yourcompany.com`. The API is implemented with API Gateway and Lambda

- A. You successfully tested the API using `curl`
- B. You implemented Javascript to call the API from a webpage on your corporate website, `www.yourcompany.co`
- C. When you access that page in your browser, you get the following error: "The same origin policy disallows reading the remote resource" How can you allow your corporate webpages to invoice the API?
- D. Disable CORS in the API Gateway.
- E. Disable CORS in the Javascript frontend.
- F. Enable CORS in the API Gateway.
- G. Enable CORS in the Javascript frontend

**Answer: D**

**NEW QUESTION 464**

Which AWS services are valid origins for an Amazon CloudFront distribution? (Select TWO.)

- A. Amazon RDS
- B. ELB Classic Load Balancer
- C. Amazon S3
- D. Amazon DynamoDB
- E. Amazon Galcier

**Answer: BC**

**NEW QUESTION 468**

Your company has separate AWS accounts for development and production. Each developer is assigned an IAM user in the development account. Developers occasionally need to access the production account to roll our changes to that environment. Your company does not allow the creation of IAM users in the production account.

What strategy will allow the development team to access the production account?

- A. Create an IAM role in the development account
- B. Allow IAM users in the development account to assume the role.
- C. Create an IAM group in the production account
- D. Grant IAM users in the development account membership in the group.
- E. Create an IAM role in the production account
- F. Allow IAM users in the development account to assume the role.
- G. Create an IAM group in the development account
- H. Grant IAM users in the development account membership in the group.

**Answer: A**

**NEW QUESTION 470**

Which services can invoke AWS Lambda functions? (Select TWO.)

- A. Amazon SNS
- B. Amazon Redshift
- C. Amazon Route53
- D. Amazon DynamoDB
- E. Elastic Load Balancing

**Answer: AD**

**NEW QUESTION 474**

You bid \$0.22 for an Amazon EC2 Spot Instance when the market price was \$0.20. For 90 minutes, the market price remained at \$0.20. Then the market price changed to \$0.25, and your instance was terminated by AWS.

What was your cost of running the instance for the entire duration?

- A. \$0.47
- B. \$0.20
- C. \$0.40
- D. \$0.22

**Answer: D**

**NEW QUESTION 476**

A stray Amazon EC2 `r3.8xlarge` instance is running in your AWS account. Before terminating it, you want to find the owner to confirm that it is not needed. Where can you find the identity that launched this instance?

- A. VPC flow logs
- B. ELB access logs
- C. CloudTrail logs
- D. Operating system logs

**Answer: C**

#### NEW QUESTION 479

You are running a web application with four Amazon EC2 instances across two Availability Zones. The instances are in an Auto Scaling group behind an ELB Classic Load Balancer. A scaling event adds one instance to the group. After the event, you notice that, although all instances are serving traffic, some instances are serving more traffic than others. Which of the following could be the problem?

- A. Cross-zone load balancing is not configured on the ELB Classic Load Balancer.
- B. Access logs are not enabled on the ELB Classic Load Balancer.
- C. A SSL/TLS certificate has not been deployed on the ELB Classic Load Balancer.
- D. Sticky bits is not enabled on the ELB Classic Load Balance

**Answer:** A

#### NEW QUESTION 480

You are running a mobile media application and are considering API Gateway for the client entry point. What benefits would this provide? (Select TWO.)

- A. Caching API responses
- B. IP blacklisting
- C. Intrusion prevention
- D. Load balancing
- E. Throttling traffic

**Answer:** AE

#### NEW QUESTION 483

Your application currently stores data on an unencrypted EBS volume. A new security policy mandates that all data must be encrypted at rest. How can you encrypt the data?

- A. Create a snapshot of the volum
- B. Create a new, encrypted volume from the snapsho
- C. Replace the volume.
- D. Stop the instanc
- E. Detach the volum
- F. Modify the EBS settings to encrypt the volum
- G. Reattach the volum
- H. Start the instance.
- I. Create a snapshot of the volum
- J. Make an encrypted copy of the snapsho
- K. Create a new volume from the new snapsho
- L. Replace the volume.
- M. Modify the EBS settings to encrypt the volum
- N. You do need to detach the volume or stop the instance.

**Answer:** A

#### NEW QUESTION 485

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

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

**Answer:** C

#### Explanation:

Amazon EC2 provides a virtual computing environments, known as an instance. After you launch an instance, AWS recommends that you check its status to confirm that it goes from the pending status to the running status, the not terminated status. The following are a few reasons why an Amazon EBS-backed instance might immediately terminate:

You've reached your volume limit.

The AMI is missing a required part. The snapshot is corrupt. References:

#### NEW QUESTION 489

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

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

**Answer:** A

#### Explanation:

Amazon ECS contains the following components:

A Cluster is a logical grouping of container instances that you can place tasks on. A Container instance is an Amazon EC2 instance that is running the Amazon ECS agent and has been registered into a cluster.

A Task definition is a description of an application that contains one or more container definitions. A Scheduler is the method used for placing tasks on container

instances. A Service is an Amazon ECS service that allows you to run and maintain a specified number of instances of a task definition simultaneously. A Task is an instantiation of a task definition that is running on a container instance. A Container is a Linux container that was created as part of a task. References:

#### NEW QUESTION 493

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

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

**Answer: C**

#### Explanation:

The user can get notifications using SNS if he has configured the notifications while creating the Auto Scaling group.

#### NEW QUESTION 495

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.

#### NEW QUESTION 500

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

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

**Answer: A**

#### Explanation:

Volume gateways provide cloud-backed storage volumes that you can mount as Internet Small Computer System Interface (iSCSI) devices from your on-premises application servers. The gateway supports the following volume configurations:

Gateway-cached volumes. You store your data in Amazon Simple Storage Service (Amazon S3) and retain a copy of frequently accessed data subsets locally.

Gateway-cached volumes offer a substantial cost savings on primary storage and minimize the need to scale your storage on-premises. You also retain low-latency access to your frequently accessed data.

Gateway-stored volumes. If you need low-latency access to your entire data set, you can configure your on-premises gateway to store all your data locally and then asynchronously back up point-in-time snapshots of this data to Amazon S3. This configuration provides durable and inexpensive offsite backups that you can recover to your local data center or Amazon EC2. For example, if you need replacement capacity for disaster recovery, you can recover the backups to Amazon EC2. References:

#### NEW QUESTION 501

You are migrating an internal server on your DC to an EC2 instance with EBS volume. Your server disk usage is around 500GB so you just copied all your data to a 2TB disk to be used with AWS Import/Export. Where will the data be imported once it arrives at Amazon?

- A. to a 2TB EBS volume
- B. to an S3 bucket with 2 objects of 1TB
- C. to an 500GB EBS volume
- D. to an S3 bucket as a 2TB snapshot

**Answer: B**

#### Explanation:

An import to Amazon EBS will have different results depending on whether the capacity of your storage device is less than or equal to 1 TB or greater than 1 TB. The maximum size of an Amazon EBS snapshot is 1 TB, so if the device image is larger than 1 TB, the image is chunked and stored on Amazon S3. The target location is determined based on the total capacity of the device, not the amount of data on the device. References:

#### NEW QUESTION 504

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

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

**Answer:** A

**Explanation:**

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

References:

**NEW QUESTION 505**

True or false: A VPC contains multiple subnets, where each subnet can span multiple Availability Zones.

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

**Answer:** C

**Explanation:**

A VPC can span several Availability Zones. In contrast, a subnet must reside within a single Availability Zone.

**NEW QUESTION 509**

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 ElasticCache to help. Which of the following statements is true in regards to ElasticCache?

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

**Answer:** D

**Explanation:**

Amazon ElastiCache is a web service that makes it easy to deploy and run Memcached or Redis protocol-compliant server nodes in the cloud. Amazon ElastiCache improves the performance of web applications by allowing you to retrieve information from a fast, managed, in-memory caching system, instead of relying entirely on slower disk-based databases. The service simplifies and offloads the management, monitoring and operation of in-memory cache environments, enabling your engineering resources to focus on developing applications.

Using Amazon ElastiCache, you can not only improve load and response times to user actions and queries, but also reduce the cost associated with scaling web applications.

**NEW QUESTION 513**

Amazon EC2 provides a \_\_\_\_\_. It is an HTTP or HTTPS request that uses the HTTP verbs GET or POST.

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

**Answer:** C

**Explanation:**

Amazon EC2 provides a Query API. These requests are HTTP or HTTPS requests that use the HTTP verbs GET or POST and a Query parameter named Action.

**NEW QUESTION 515**

An organization has three separate AWS accounts, one each for development, testing, and production. The organization wants the testing team to have access to certain AWS resources in the production account. How can the organization achieve this?

- A. It is not possible to access resources of one account with another account.
- B. Create the IAM roles with cross account access.
- C. Create the IAM user in a test account, and allow it access to the production environment with the IAM policy.
- D. Create the IAM users with cross account access

**Answer:** B

**Explanation:**

An organization has multiple AWS accounts to isolate a development environment from a testing or production environment. At times the users from one account need to access resources in the other account, such as promoting an update from the development environment to the production environment. In this case the IAM role with cross account access will provide a solution. Cross account access lets one account share access to their resources with users in the other AWS accounts.

#### NEW QUESTION 517

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.

#### NEW QUESTION 522

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

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

**Answer: B**

#### Explanation:

By default, all AWS accounts are limited to 5 Elastic IP addresses per region for each AWS account, because public (IPv4) Internet addresses are a scarce public resource. AWS strongly encourages you to use an EIP primarily for load balancing use cases, and use DNS hostnames for all other inter-node communication.

If you feel your architecture warrants additional EIPs, you would need to complete the Amazon EC2 Elastic IP Address Request Form and give reasons as to your need for additional addresses. References:

#### NEW QUESTION 526

In Amazon EC2, partial instance-hours are billed \_\_\_\_\_.

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

**Answer: D**

#### Explanation:

Partial instance-hours are billed to the next hour. References:

#### NEW QUESTION 529

You are setting up a VPC and you need to set up a public subnet within that VPC. Which following requirement must be met for this subnet to be considered a public subnet?

- A. Subnet's traffic is not routed to an internet gateway but has its traffic routed to a virtual private gateway.
- B. Subnet's traffic is routed to an internet gateway.
- C. Subnet's traffic is not routed to an internet gateway.
- D. None of these answers can be considered a public subne

**Answer: B**

#### Explanation:

A virtual private cloud (VPC) is a virtual network dedicated to your AWS account. It is logically isolated from other virtual networks in the AWS cloud. You can launch your AWS resources, such as Amazon EC2 instances, into your VPC. You can configure your VPC: you can select its IP address range, create subnets, and configure route tables, network gateways, and security settings. A subnet is a range of IP addresses in your VPC. You can launch AWS resources into a subnet that you select. Use a public subnet for resources that must be connected to the Internet, and a private subnet for resources that won't be connected to the Internet.

If a subnet's traffic is routed to an internet gateway, the subnet is known as a public subnet. If a subnet doesn't have a route to the internet gateway, the subnet is known as a private subnet. If a subnet doesn't have a route to the internet gateway, but has its traffic routed to a virtual private gateway, the subnet is known as a VPN-only subnet.

#### NEW QUESTION 531

.....

## **Thank You for Trying Our Product**

### **We offer two products:**

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

2nd - Questions and Answers in PDF Format

### **SAA-C01 Practice Exam Features:**

- \* SAA-C01 Questions and Answers Updated Frequently
- \* SAA-C01 Practice Questions Verified by Expert Senior Certified Staff
- \* SAA-C01 Most Realistic Questions that Guarantee you a Pass on Your First Try
- \* SAA-C01 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The SAA-C01 Practice Test Here](#)**