

# Exam Questions AWS-Solution-Architect-Associate

Amazon AWS Certified Solutions Architect - Associate

<https://www.2passeasy.com/dumps/AWS-Solution-Architect-Associate/>



#### NEW QUESTION 1

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

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

**Answer:** A

#### Explanation:

Amazon ECS contains the following components:

A Cluster is a logical grouping of container instances that you can place tasks on.

A Container instance is an Amazon EC2 instance that is running the Amazon ECS agent and has been registered into a cluster.

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

A Service is an Amazon ECS service that allows you to run and maintain a specified number of instances of a task definition simultaneously.

A Task is an instantiation of a task definition that is running on a container instance. A Container is a Linux container that was created as part of a task.

Reference: <http://docs.aws.amazon.com/AmazonECS/latest/developerguide/Welcome.html>

#### NEW QUESTION 2

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

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

**Answer:** C

#### Explanation:

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

Reference: <http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/GettingStartedTutorial.html>

#### NEW QUESTION 3

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

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

**Answer:** A

#### Explanation:

Amazon allows you to make backups of the data stored in your EBS volumes through snapshots that can later be used to create a new EBS volume.

Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/Storage.html>

#### NEW QUESTION 4

To specify a resource in a policy statement, in Amazon EC2, can you use its Amazon Resource Name (ARN)?

- A. Yes, you can.
- B. No, you can't because EC2 is not related to ARN.
- C. No, you can't because you can't specify a particular Amazon EC2 resource in an IAM policy.
- D. Yes, you can but only for the resources that are not affected by the action

**Answer:** A

#### Explanation:

Some Amazon EC2 API actions allow you to include specific resources in your policy that can be created or modified by the action. To specify a resource in the statement, you need to use its Amazon Resource Name (ARN).

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-ug.pdf>

#### NEW QUESTION 5

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 off-site backups that you can recover to your local data center or Amazon EC2. For example, if you need replacement capacity for disaster recovery, you can recover the backups to Amazon EC2.

Reference: <http://docs.aws.amazon.com/storagegateway/latest/userguide/volume-gateway.html>

**NEW QUESTION 6**

After setting up a Virtual Private Cloud (VPC) network, a more experienced cloud engineer suggests that to achieve low network latency and high network throughput you should look into setting up a placement group. You know nothing about this, but begin to do some research about it and are especially curious about its limitations. Which of the below statements is wrong in describing the limitations of a placement group?

- A. Although launching multiple instance types into a placement group is possible, this reduces the likelihood that the required capacity will be available for your launch to succeed.
- B. A placement group can span multiple Availability Zones.
- C. You can't move an existing instance into a placement group.
- D. A placement group can span peered VPCs

**Answer: B**

**Explanation:**

A placement group is a logical grouping of instances within a single Availability Zone. Using placement groups enables applications to participate in a low-latency, 10 Gbps network. Placement groups are recommended for applications that benefit from low network latency, high network throughput, or both. To provide the lowest latency, and the highest packet-per-second network performance for your placement group, choose an instance type that supports enhanced networking. Placement groups have the following limitations:

The name you specify for a placement group a name must be unique within your AWS account. A placement group can't span multiple Availability Zones.

Although launching multiple instance types into a placement group is possible, this reduces the likelihood that the required capacity will be available for your launch to succeed. We recommend using the same instance type for all instances in a placement group.

You can't merge placement groups. Instead, you must terminate the instances in one placement group, and then relaunch those instances into the other placement group.

A placement group can span peered VPCs; however, you will not get full-bisection bandwidth between instances in peered VPCs. For more information about VPC peering connections, see VPC Peering in the Amazon VPC User Guide.

You can't move an existing instance into a placement group. You can create an AM from your existing instance, and then launch a new instance from the AMI into a placement group.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>

**NEW QUESTION 7**

What is a placement group in Amazon EC2?

- A. It is a group of EC2 instances within a single Availability Zone.
- B. It the edge location of your web content.
- C. It is the AWS region where you run the EC2 instance of your web content.
- D. It is a group used to span multiple Availability Zone

**Answer: A**

**Explanation:**

A placement group is a logical grouping of instances within a single Availability Zone. Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>

**NEW QUESTION 8**

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.

Reference: <http://docs.aws.amazon.com/AWSImportExport/latest/DG/Concepts.html>

**NEW QUESTION 9**

An edge location refers to which Amazon Web Service?

- A. An edge location is referred to the network configured within a Zone or Region
- B. An edge location is an AWS Region

- C. An edge location is the location of the data center used for Amazon CloudFront.
- D. An edge location is a Zone within an AWS Region

**Answer:** C

**Explanation:**

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

**NEW QUESTION 10**

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 10**

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

- A. Amazon SES
- B. Amazon CloudSearch
- C. Amazon SWF
- D. Amazon AppStream

**Answer:** A

**Explanation:**

File change notifications can be sent via email to users following the resource with Amazon Simple Email Service (Amazon SES), an easy-to-use, cost-effective email solution. Reference: [http://media.amazonwebservices.com/architecturecenter/AWS\\_ac\\_ra\\_filesync\\_08.pdf](http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_filesync_08.pdf)

**NEW QUESTION 15**

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

- A. Amazon CloudFront
- B. AWS Direct Connect
- C. AWS CloudHSM
- D. AWS VPN CloudHub

**Answer:** D

**Explanation:**

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

**NEW QUESTION 19**

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

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

**Answer:** B

**Explanation:**

Key pairs consist of a public and private key, where you use the private key to create a digital signature, and then AWS uses the corresponding public key to validate the signature. Key pairs are used only for Amazon EC2 and Amazon CloudFront.

Reference: <http://docs.aws.amazon.com/general/latest/gr/aws-sec-cred-types.html>

**NEW QUESTION 22**

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.

Reference: [http://media.amazonwebservices.com/AWS\\_Security\\_Best\\_Practices.pdf](http://media.amazonwebservices.com/AWS_Security_Best_Practices.pdf)

**NEW QUESTION 24**

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

- A. Oracle export/import utilities
- B. Oracle SQL Developer
- C. Oracle Data Pump
- D. DBMS\_FILE\_TRANSFER

**Answer:** C

**Explanation:**

How you import data into an Amazon RDS DB instance depends on the amount of data you have and the number and variety of database objects in your database.

For example, you can use Oracle SQL Developer to import a simple, 20 MB database; you want to use Oracle Data Pump to import complex databases or databases that are several hundred megabytes or several terabytes in size.

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

**NEW QUESTION 29**

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

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

**Answer:** C

**Explanation:**

AWS Import/Export supports: Import to Amazon S3

Export from Amazon S3 Import to Amazon EBS Import to Amazon Glacier

AWS Import/Export does not currently support export from Amazon EBS or Amazon Glacier. Reference:

<https://docs.aws.amazon.com/AWSImportExport/latest/DG/whatisdisk.html>

**NEW QUESTION 31**

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

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

**Answer:** B

**Explanation:**

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

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

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-ip-addresses-eip.html#using-instance-addressing-limit>

**NEW QUESTION 33**

In Amazon EC2, partial instance-hours are billed .

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

**Answer:** D

**Explanation:**

Partial instance-hours are billed to the next hour. Reference: <http://aws.amazon.com/ec2/faqs/>

**NEW QUESTION 37**

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

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

**Answer:** B

**Explanation:**

A virtual private cloud (VPC) is a virtual network dedicated to your AWS account. It is logically isolated from other virtual networks in the AWS cloud. You can launch your AWS resources, such as Amazon EC2 instances, into your VPC. You can configure your VPC: you can select its IP address range, create subnets, and configure route tables, network gateways, and security settings.

A subnet is a range of IP addresses in your VPC. You can launch AWS resources into a subnet that you select. Use a public subnet for resources that must be connected to the internet, and a private subnet for resources that won't be connected to the Internet.

If a subnet's traffic is routed to an internet gateway, the subnet is known as a public subnet.

If a subnet doesn't have a route to the internet gateway, the subnet is known as a private subnet.

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

Reference: [http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC\\_Subnets.html](http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html)

**NEW QUESTION 39**

Can you specify the security group that you created for a VPC when you launch an instance in EC2-Classic?

- A. No, you can specify the security group created for EC2-Classic when you launch a VPC instance.
- B. No
- C. Yes
- D. No, you can specify the security group created for EC2-Classic to a non-VPC based instance onl

**Answer:** B

**Explanation:**

If you're using EC2-Classic, you must use security groups created specifically for EC2-Classic. When you launch an instance in EC2-Classic, you must specify a security group in the same region as the instance. You can't specify a security group that you created for a VPC when you launch an instance in EC2-Classic.

Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html#ec2-classic-security-groups>

**NEW QUESTION 42**

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

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

**Answer:** C

**Explanation:**

Amazon Redshift achieves efficient storage and optimum query performance through a combination of massively parallel processing, columnar data storage, and very efficient, targeted data compression encoding schemes.

Columnar storage for database tables is an important factor in optimizing analytic query performance because it drastically reduces the overall disk I/O requirements and reduces the amount of data you need to load from disk.

Reference: [http://docs.aws.amazon.com/redshift/latest/dg/c\\_columnar\\_storage\\_disk\\_mem\\_mgmt.html](http://docs.aws.amazon.com/redshift/latest/dg/c_columnar_storage_disk_mem_mgmt.html)

**NEW QUESTION 44**

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

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

**Answer:** C

**Explanation:**

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

If your I/O latency is higher than you require, check your average queue length to make sure that your application is not trying to drive more IOPS than you have provisioned. You can maintain high IOPS while keeping latency down by maintaining a low average queue length (which is achieved by provisioning more IOPS for your volume).

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-workload-demand.html>

**NEW QUESTION 48**

You have been given a scope to deploy some AWS infrastructure for a large organisation. The requirements are that you will have a lot of EC2 instances but may need to add more when the average utilization of your Amazon EC2 fleet is high and conversely remove them when CPU utilization is low. Which AWS services would be best to use to accomplish this?

- A. Auto Scaling, Amazon CloudWatch and AWS Elastic Beanstalk
- B. Auto Scaling, Amazon CloudWatch and Elastic Load Balancing.
- C. Amazon CloudFront, Amazon CloudWatch and Elastic Load Balancing.
- D. AWS Elastic Beanstalk , Amazon CloudWatch and Elastic Load Balancing

**Answer:** B

**Explanation:**

Auto Scaling enables you to follow the demand curve for your applications closely, reducing the need to manually provision Amazon EC2 capacity in advance. For example, you can set a condition to add new

Amazon EC2 instances in increments to the Auto Scaling group when the average utilization of your Amazon EC2 fleet is high; and similarly, you can set a condition to remove instances in the same increments when CPU utilization is low. If you have predictable load changes, you can set a schedule through Auto Scaling to plan your scaling activities. You can use Amazon CloudWatch to send alarms to trigger scaling activities and Elastic Load Balancing to help distribute traffic to your instances within Auto Scaling groups. Auto Scaling enables you to run your Amazon EC2 fleet at optimal utilization. Reference:

<http://aws.amazon.com/autoscaling/>

**NEW QUESTION 53**

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

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

**Answer:** A

**Explanation:**

Several factors can affect the performance of Amazon EBS volumes, such as instance configuration, I/O characteristics, workload demand, and storage configuration.

Frequent snapshots provide a higher level of data durability, but they may slightly degrade the performance of your application while the snapshot is in progress. This trade off becomes critical when you have data that changes rapidly. Whenever possible, plan for snapshots to occur during off-peak times in order to minimize workload impact.

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

**NEW QUESTION 54**

You've created your first load balancer and have registered your EC2 instances with the load balancer. Elastic Load Balancing routinely performs health checks on all the registered EC2 instances and automatically distributes all incoming requests to the DNS name of your load balancer across your registered, healthy EC2 instances. By default, the load balancer uses the \_ protocol for checking the health of your instances.

- A. HTTPS
- B. HTTP
- C. ICMP
- D. IPv6

**Answer:** B

**Explanation:**

In Elastic Load Balancing a health configuration uses information such as protocol, ping port, ping path (URL), response timeout period, and health check interval to determine the health state of the instances registered with the load balancer.

Currently, HTTP on port 80 is the default health check. Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/TerminologyandKeyConcepts.html>

**NEW QUESTION 58**

In Amazon EC2 Container Service, are other container types supported?

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 59**

is a fast, filexible, fully managed push messaging service.

- A. Amazon SNS
- B. Amazon SES
- C. Amazon SQS
- D. Amazon FPS

**Answer:** A

**Explanation:**

Amazon Simple Notification Service (Amazon SNS) is a fast, filexible, fully managed push messaging service. Amazon SNS makes it simple and cost-effective to push to mobile devices such as iPhone, iPad, Android, Kindle Fire, and internet connected smart devices, as well as pushing to other distributed services.  
Reference: [http://aws.amazon.com/sns/?nc1=h\\_I2\\_as](http://aws.amazon.com/sns/?nc1=h_I2_as)

**NEW QUESTION 62**

In Amazon EC2, while sharing an Amazon EBS snapshot, can the snapshots with AWS IV|arketplace product codes be public?

- A. Yes, but only for US-based providers.
- B. Yes, they can be public.
- C. No, they cannot be made public.
- D. Yes, they are automatically made public by the syste

**Answer:** C

**Explanation:**

Snapshots with AWS Marketplace product codes can't be made public. Reference:  
<http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/ebs-modifying-snapshot-permissions.html>

**NEW QUESTION 67**

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 70**

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 74**

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

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

**Answer:** A

**Explanation:**

Amazon RDS supports DB instances running several versions of PostgreSQL. Currently we support PostgreSQL versions 9.3.1, 9.3.2, and 9.3.3. You can create DB instances and DB snapshots, point-in-time restores and backups. DB instances running PostgreSQL support Multi-AZ deployments, Provisioned IOPS, and can be created inside a VPC. You can also use SSL to connect to a DB instance running PostgreSQL.

You can use any standard SQL client application to run commands for the instance from your client computer. Such applications include pgAdmin, a popular Open Source administration and development tool for PostgreSQL, or psql, a command line utility that is part of a PostgreSQL installation. In order to deliver a managed service experience, Amazon RDS does not provide host access to DB instances, and it restricts access to certain system procedures and tables that require advanced privileges. Amazon RDS supports access to databases on a DB instance using any standard SQL client application. Amazon RDS does not allow direct host access to a DB instance via Telnet or Secure Shell (SSH).

Reference: [http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP\\_PostgreSQL.html](http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_PostgreSQL.html)

**NEW QUESTION 79**

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

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

**Answer:** A

**Explanation:**

Amazon DynamoDB is ideal for database applications that require very low latency and predictable performance at any scale but don't need complex querying capabilities like joins or transactions. Amazon DynamoDB is a fully-managed NoSQL database service that offers high performance, predictable throughput and low cost. It is easy to set up, operate, and scale.

With Amazon DynamoDB, you can start small, specify the throughput and storage you need, and easily scale your capacity requirements on the fly. Amazon DynamoDB automatically partitions data over a number of servers to meet your request capacity. In addition, DynamoDB automatically replicates your data synchronously across multiple Availability Zones within an AWS Region to ensure high-availability and data durability.

Reference: [https://aws.amazon.com/running\\_databases/#dynamodb\\_anchor](https://aws.amazon.com/running_databases/#dynamodb_anchor)

**NEW QUESTION 84**

Once again your customers are concerned about the security of their sensitive data and with their latest enquiry ask about what happens to old storage devices on AWS. What would be the best answer to this QUESTION ?

- A. AWS reformats the disks and uses them again.
- B. AWS uses the techniques detailed in DoD 5220.22-M to destroy data as part of the decommissioning process.
- C. AWS uses their own proprietary software to destroy data as part of the decommissioning process.
- D. AWS uses a 3rd party security organization to destroy data as part of the decommissioning process

**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.

Reference: <http://d0.awsstatic.com/whitepapers/Security/AWS%20Security%20Whitepaper.pdf>

**NEW QUESTION 85**

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

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

**Answer:** D

**Explanation:**

Amazon RDS provides high availability and failover support for DB instances using Multi-AZ deployments. Multi-AZ deployments for Oracle, PostgreSQL, MySQL, and MariaDB DB instances use Amazon technology, while SQL Server DB instances use SQL Server Mirroring.

Reference: <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.IV|u|tiAZ.html>

**NEW QUESTION 87**

What does Amazon DynamoDB provide?

- A. A predictable and scalable MySQL database
- B. A fast and reliable PL/SQL database cluster
- C. A standalone Cassandra database, managed by Amazon Web Services
- D. A fast, highly scalable managed NoSQL database service

**Answer:** D

**Explanation:**

Amazon DynamoDB is a managed NoSQL database service offered by Amazon. It automatically manages tasks like scalability for you while it provides high availability and durability for your data, allowing you to concentrate in other aspects of your application.

Reference: check link - [https://aws.amazon.com/running\\_databases/](https://aws.amazon.com/running_databases/)

#### NEW QUESTION 88

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

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

**Answer: D**

#### Explanation:

To determine your instance's public IP address from within the instance, you can use instance metadata. Use the following command to access the public IP address: For Linux use, \$ curl

<http://169.254.169.254/latest/meta-data/public-ipv4>, and for Windows use, \$ wget <http://169.254.169.254/latest/meta-data/public-ipv4>.

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

#### NEW QUESTION 91

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

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

**Answer: D**

#### Explanation:

Several factors can affect the performance of Amazon EBS volumes, such as instance configuration, I/O characteristics, workload demand, and storage configuration.

IOPS are input/output operations per second. Amazon EBS measures each I/O operation per second

(that is 256 KB or smaller) as one IOPS. I/O operations that are larger than 256 KB are counted in 256 KB capacity units.

For example, a 1,024 KB I/O operation would count as 4 IOPS.

When you provision a 4,000 IOPS volume and attach it to an EBS-optimized instance that can provide the necessary bandwidth, you can transfer up to 4,000 chunks of data per second (provided that the I/O does not exceed the 128 MB/s per volume throughput limit of General Purpose (SSD) and Provisioned IOPS (SSD) volumes).

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

#### NEW QUESTION 94

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

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

**Answer: A**

#### Explanation:

You can write an AWS CloudFormation template (a JSON-formatted document) in a text editor or pick an existing template. The template describes the resources you want and their settings. For example,

suppose you want to create an Amazon EC2. Your template can declare an instance Amazon EC2 and describe its properties, as shown in the following example:

```
{
  "AWSTemplateFormatVersion" : "2010-09-09",
  "Description" : "A simple Amazon EC2 instance", "Resources" : {
    "MyEC2Instance" : {
      "Type" : "AWS::EC2::Instance", "Properties" : {
        "ImageId" : "ami-2f726546", "InstanceType" : "t1.micro"
      }
    }
  }
}
```

Reference:

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-what-is-how-does-it-work.html>

#### NEW QUESTION 98

Which of the following is NOT a characteristic of Amazon Elastic Compute Cloud (Amazon EC2)?

- A. It can be used to launch as many or as few virtual servers as you need.
- B. It increases the need to forecast traffic by providing dynamic IP addresses for static cloud computing.
- C. It eliminates your need to invest in hardware up front, so you can develop and deploy applications faster.
- D. It offers scalable computing capacity in the Amazon Web Services (AWS) cloud.

**Answer:** B

**Explanation:**

Amazon Elastic Compute Cloud (Amazon EC2) provides scalable computing capacity in the Amazon Web Services (AWS) cloud. Using Amazon EC2 eliminates your need to invest in hardware up front, so you can develop and deploy applications faster. You can use Amazon EC2 to launch as many or as few virtual servers as you need, configure security and networking, and manage storage. Amazon EC2 enables you to scale up or down to handle changes in requirements or spikes in popularity, reducing your need to forecast traffic.  
Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/concepts.html>

**NEW QUESTION 101**

You are setting up your first Amazon Virtual Private Cloud (Amazon VPC) so you decide to use the VPC wizard in the AWS console to help make it easier for you. Which of the following statements is correct regarding instances that you launch into a default subnet via the VPC wizard?

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 105**

Does AWS CloudFormation support Amazon EC2 tagging?

- A. Yes, AWS CloudFormation supports Amazon EC2 tagging
- B. No, CloudFormation doesn't support any tagging
- C. No, it doesn't support Amazon EC2 tagging.
- D. It depends if the Amazon EC2 tagging has been defined in the template

**Answer:** A

**Explanation:**

In AWS CloudFormation, Amazon EC2 resources that support the tagging feature can also be tagged in an AWS template. The tag values can refer to template parameters, other resource names, resource attribute values (e.g. addresses), or values computed by simple functions (e.g., a concatenated list of strings).  
Reference: <http://aws.amazon.com/cloudformation/faqs/>

**NEW QUESTION 106**

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 111**

Select a true statement about Amazon EC2 Security Groups (EC2-Classic).

- A. After you launch an instance in EC2-Classic, you can't change its security groups.
- B. After you launch an instance in EC2-Classic, you can change its security groups only once.
- C. After you launch an instance in EC2-Classic, you can only add rules to a security group.
- D. After you launch an instance in EC2-Classic, you cannot add or remove rules from a security group

**Answer:** A

**Explanation:**

After you launch an instance in EC2-Classic, you can't change its security groups. However, you can add rules to or remove rules from a security group, and those changes are automatically applied to all instances that are associated with the security group.

Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/using-network-security.html>

#### NEW QUESTION 114

A user is observing the EC2 CPU utilization metric on CloudWatch. The user has observed some interesting patterns while filtering over the 1 week period for a particular hour. The user wants to zoom that data point to a more granular period. How can the user do that easily with CloudWatch?

- A. The user can zoom a particular period by selecting that period with the mouse and then releasing the mouse
- B. The user can zoom a particular period by specifying the aggregation data for that period
- C. The user can zoom a particular period by double clicking on that period with the mouse
- D. The user can zoom a particular period by specifying the period in the Time Range

**Answer:** A

#### Explanation:

Amazon CloudWatch provides the functionality to graph the metric data generated either by the AWS services or the custom metric to make it easier for the user to analyse. The AWS CloudWatch console provides the option to change the granularity of a graph and zoom in to see data over a shorter time period. To zoom, the user has to click in the graph details pane, drag on the graph area for selection, and then release the mouse button.

Reference: [http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/zoom\\_in\\_on\\_graph.html](http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/zoom_in_on_graph.html)

#### NEW QUESTION 119

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

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

**Answer:** D

#### Explanation:

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

Reference: [http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/device\\_naming.html](http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/device_naming.html)

#### NEW QUESTION 124

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

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

**Answer:** D

#### Explanation:

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

Reference: [http://aws.amazon.com/dynamodb/faqs/#security\\_anchor](http://aws.amazon.com/dynamodb/faqs/#security_anchor)

#### NEW QUESTION 125

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

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

**Answer:** A

#### Explanation:

Amazon CloudWatch integrates with AWS Identity and Access Management (IAM) so that you can specify which CloudWatch actions a user in your AWS Account can perform. For example, you could create an IAM policy that gives only certain users in your organization permission to use GetMetricStatistics. They could then use the action to retrieve data about your cloud resources.

You can't use IAM to control access to CloudWatch data for specific resources. For example, you can't give a user access to CloudWatch data for only a specific set of instances or a specific LoadBalancer. Permissions granted using IAM cover all the cloud resources you use with CloudWatch. In addition, you can't use IAM roles with the Amazon CloudWatch command line tools.

Using Amazon CloudWatch with IAM doesn't change how you use CloudWatch. There are no changes to CloudWatch actions, and no new CloudWatch actions related to users and access control.

Reference: <http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/UsingIAM.html>

#### NEW QUESTION 127

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

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

**Answer:** D

**Explanation:**

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

Reference: [http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER\\_WorkingWithSecurityGroups.html](http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithSecurityGroups.html)

**NEW QUESTION 130**

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

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

**Answer:** B

**Explanation:**

Amazon Simple Notification Service (Amazon SNS) is a fast, flexible, and fully managed push messaging service. Amazon SNS can be used to make push notifications to mobile devices. Amazon SNS can

deliver notifications by SMS text message or email to the Amazon Simple Queue Service (SQS) queues or to any HTTP endpoint. In this case user can use the SNS apis to send SMS.

Reference: <http://aws.amazon.com/sns/>

**NEW QUESTION 131**

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

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

**Answer:** B

**Explanation:**

The AWS Storage Gateway is a service connecting an on-premises software appliance with cloud-based storage to provide seamless and secure integration between an organization's on-premises IT environment and AWS's storage infrastructure.

AWS Storage Gateway billing is as follows. Volume storage usage (per GB per month):

You are billed for the Cached volume data you store in Amazon S3. You are only billed for volume capacity you use, not for the size of the volume you create.

Snapshot Storage usage (per GB per month): You are billed for the snapshots your gateway stores in Amazon S3. These snapshots are stored and billed as Amazon EBS snapshots. Snapshots are incremental backups, reducing your storage charges. When taking a new snapshot, only the data that has changed since your last snapshot is stored.

Virtual Tape Library usage (per GB per month):

You are billed for the virtual tape data you store in Amazon S3. You are only billed for the portion of virtual tape capacity that you use, not for the size of the virtual tape.

Virtual Tape Shelf usage (per GB per month):

You are billed for the virtual tape data you store in Amazon Glacier. You are only billed for the portion of virtual tape capacity that you use, not for the size of the virtual tape.

Reference: <https://aws.amazon.com/storagegateway/faqs/>

**NEW QUESTION 134**

You are configuring a new VPC for one of your clients for a cloud migration project, and only a public VPN will be in place. After you created your VPC, you created a new subnet, a new internet gateway, and attached your internet gateway to your VPC. When you launched your first instance into your VPC, you realized that you aren't able to connect to the instance, even if it is configured with an elastic IP. What should be done to access the instance?

- A. A route should be created as 0.0.0.0/0 and your internet gateway as target.
- B. Attach another ENI to the instance and connect via new ENI.
- C. A NAT instance should be created and all traffic should be forwarded to NAT instance.
- D. A NACL should be created that allows all outbound traffic

**Answer:** A

**Explanation:**

All traffic should be routed via Internet Gateway. So, a route should be created with 0.0.0.0/0 as a source, and your Internet Gateway as your target.

Reference: [http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC\\_Scenario1.html](http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario1.html)

**NEW QUESTION 138**

You receive a bill from AWS but are confused because you see you are incurring different costs for the exact same storage size in different regions on Amazon S3. You ask AWS why this is so. What response would you expect to receive from AWS?

- A. We charge less in different time zones.

- B. We charge less where our costs are less.
- C. This will balance out next bill.
- D. It must be a mistake

**Answer:** B

**Explanation:**

Amazon S3 is storage for the internet. It's a simple storage service that offers software developers a highly-scalable, reliable, and low-latency data storage infrastructure at very low costs.

AWS charges less where their costs are less.

For example, their costs are lower in the US Standard Region than in the US West (Northern California) Region.

Reference: <https://aws.amazon.com/s3/faqs/>

**NEW QUESTION 141**

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

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

**Answer:** B

**Explanation:**

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

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

**NEW QUESTION 142**

What is the default maximum number of Access Keys per user?

- A. 10
- B. 15
- C. 2
- D. 20

**Answer:** C

**Explanation:**

The default maximum number of Access Keys per user is 2.

Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/LimitationsOnEntities.html>

**NEW QUESTION 147**

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

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

**Answer:** B

**Explanation:**

In Amazon VPC, you can assign any Private IP address to your instance as long as it is: Part of the associated subnet's IP address range

Not reserved by Amazon for IP networking purposes Not currently assigned to another interface Reference: <http://aws.amazon.com/vpc/faqs/>

**NEW QUESTION 150**

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

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

**Answer:** C

**Explanation:**

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

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-ec2-config.html>

#### NEW QUESTION 155

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

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

**Answer:** D

#### Explanation:

Access to Amazon S3 from within Amazon EC2 in the same region is fast. In this aspect, though the client base is Singapore, the application is being hosted in the US West-1 region. Thus, it is recommended that S3 objects be stored in the US-West-1 region.

Reference: [http://media.amazonwebservices.com/AWS\\_Storage\\_Options.pdf](http://media.amazonwebservices.com/AWS_Storage_Options.pdf)

#### NEW QUESTION 159

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

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

**Answer:** C

#### Explanation:

You can assign tags only to resources that already exist. You can't terminate, stop, or delete a resource based solely on its tags; you must specify the resource identifier.

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

#### NEW QUESTION 161

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

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

**Answer:** C

#### Explanation:

IAM users need passwords in order to access the AWS Management Console. (They do not need passwords if they will access AWS resources programmatically by using the CLI, AWS SDKs, or the APIs.)

You can use a password policy to do these things: Set a minimum password length.

Require specific character types, including uppercase letters, lowercase letters, numbers, and non-alphanumeric characters. Be sure to remind your users that passwords are case sensitive. Allow all IAM users to change their own passwords.

Require IAM users to change their password after a specified period of time (enable password expiration). Prevent IAM users from reusing previous passwords.

Force IAM users to contact an account administrator when the user has allowed his or her password to expire.

Reference: [http://docs.aws.amazon.com/IAM/latest/UserGuide/Using\\_ManagingPasswordPolicies.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_ManagingPasswordPolicies.html)

#### NEW QUESTION 162

A user has created a CloudFormation stack. The stack creates AWS services, such as EC2 instances, ELB, AutoScaling, and RDS. While creating the stack it created EC2, ELB and AutoScaling but failed to create RDS. What will CloudFormation do in this scenario?

- A. Rollback all the changes and terminate all the created services
- B. It will wait for the user's input about the error and correct the mistake after the input
- C. CloudFormation can never throw an error after launching a few services since it verifies all the steps before launching
- D. It will warn the user about the error and ask the user to manually create RDS

**Answer:** A

#### Explanation:

AWS CloudFormation is an application management tool which provides application modeling, deployment, configuration, management and related activities. The AWS CloudFormation stack is a collection of AWS resources which are created and managed as a single unit when AWS CloudFormation instantiates a template. If any of the services fails to launch, CloudFormation will rollback all the changes and terminate or delete all the created services.

Reference: <http://aws.amazon.com/cloudformation/faqs/>

#### NEW QUESTION 167

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

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

**Answer:** B

**Explanation:**

AWS Direct Connect is a network service that provides an alternative to using the internet to utilize AWS cloud services.

Using AWS Direct Connect, data that would have previously been transported over the Internet can now be delivered through a private network connection between AWS and your datacenter or corporate network.

1Gbps and 10Gbps ports are available. Speeds of 50Mbps, 100Mbps, 200Mbps, 300Mbps, 400Mbps, and 500Mbps can be ordered from any APN partners supporting AWS Direct Connect.

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

**NEW QUESTION 168**

After setting up some EC2 instances you now need to set up a monitoring solution to keep track of these instances and to send you an email when the CPU hits a certain threshold. Which statement below best describes what thresholds you can set to trigger a CloudWatch Alarm?

- A. Set a target value and choose whether the alarm will trigger when the value is greater than (>), greater than or equal to (>=), less than (<), or less than or equal to (<=) that value.
- B. Thresholds need to be set in IAM not CloudWatch
- C. Only default thresholds can be set you can't choose your own thresholds.
- D. Set a target value and choose whether the alarm will trigger when the value hits this threshold

**Answer:** A

**Explanation:**

Amazon CloudWatch is a monitoring service for AWS cloud resources and the applications you run on AWS. You can use Amazon CloudWatch to collect and track metrics, collect and monitor log files, and set alarms.

When you create an alarm, you first choose the Amazon CloudWatch metric you want it to monitor. Next, you choose the evaluation period (e.g., five minutes or one hour) and a statistical value to measure (e.g., Average or Maximum).

To set a threshold, set a target value and choose whether the alarm will trigger when the value is greater than (>), greater than or equal to (>=), less than (<), or less than or equal to (<=) that value.

Reference: <http://aws.amazon.com/cloudwatch/faqs/>

**NEW QUESTION 172**

Which of the following would you use to list your AWS Import/Export jobs?

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

**Answer:** C

**Explanation:**

You can list AWS Import/Export jobs with the ListJobs command using the command line client or REST API.

Reference: <http://docs.aws.amazon.com/AWSImportExport/latest/DG/ListingYourJobs.html>

**NEW QUESTION 174**

A favored client needs you to quickly deploy a database that is a relational database service with minimal administration as he wants to spend the least amount of time administering it. Which database would be the best option?

- A. Amazon SimpleDB
- B. Your choice of relational AMs on Amazon EC2 and EBS.
- C. Amazon RDS
- D. Amazon Redshift

**Answer:** C

**Explanation:**

Amazon Relational Database Service (Amazon RDS) is a web service that makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity while managing time-consuming database administration tasks, freeing you up to focus on your applications and business.

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

Reference: [https://aws.amazon.com/running\\_databases/#rds\\_anchor](https://aws.amazon.com/running_databases/#rds_anchor)

**NEW QUESTION 175**

How many types of block devices does Amazon EC2 support?

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

**Answer:** C

**Explanation:**

Amazon EC2 supports 2 types of block devices. Reference:

<http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/block-device-mapping-concepts.html>

#### NEW QUESTION 180

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

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

**Answer:** A

#### Explanation:

Amazon Elastic Load Balancing is used to manage traffic on a fleet of Amazon EC2 instances, distributing traffic to instances across all availability zones within a region. Elastic Load Balancing has all the advantages of an on-premises load balancer, plus several security benefits:

Takes over the encryption and decryption work from the Amazon EC2 instances and manages it centrally on the load balancer

Offers clients a single point of contact, and can also serve as the first line of defense against attacks on your network

When used in an Amazon VPC, supports creation and management of security groups associated with your Elastic Load Balancing to provide additional networking and security options

Supports end-to-end traffic encryption using TLS (previously SSL) on those networks that use secure HTTP (HTTPS) connections. When TLS is used, the TLS server certificate used to terminate client connections can be managed centrally on the load balancer, rather than on every individual instance. Reference:

<http://d0.awsstatic.com/whitepapers/Security/AWS%20Security%20Whitepaper.pdf>

#### NEW QUESTION 181

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

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

**Answer:** C

#### Explanation:

You can use the AWS Policy Generator to create a bucket policy for your Amazon S3 bucket. You can then use the generated document to set your bucket policy by using the Amazon S3 console, by a number of third-party tools, or via your application.

You use an Amazon S3 bucket policy that specifies a wildcard (\*) in the principal element, which means anyone can access the bucket. With anonymous access, anyone (including users without an AWS account) will be able to access the bucket.

Reference: <http://docs.aws.amazon.com/IAM/latest/UserGuide/iam-troubleshooting.htm#d0e20565>

#### NEW QUESTION 183

You have been asked to build AWS infrastructure for disaster recovery for your local applications and within that you should use an AWS Storage Gateway as part of the solution. Which of the following best describes the function of an AWS Storage Gateway?

- A. Accelerates transferring large amounts of data between the AWS cloud and portable storage devices .
- B. A web service that speeds up distribution of your static and dynamic web content.
- C. Connects an on-premises software appliance with cloud-based storage to provide seamless and secure integration between your on-premises IT environment and AWS's storage infrastructure.
- D. Is a storage service optimized for infrequently used data, or "cold data."

**Answer:** C

#### Explanation:

AWS Storage Gateway connects an on-premises software appliance with cloud-based storage to provide seamless integration with data security features between your on-premises IT environment and the Amazon Web Services (AWS) storage infrastructure. You can use the service to store data in the AWS cloud for scalable and cost-effective storage that helps maintain data security. AWS Storage Gateway offers both volume-based and tape-based storage solutions:

Volume gateways Gateway-cached volumes Gateway-stored volumes

Gateway-virtual tape library (VTL)

Reference:

[http://media.amazonwebservices.com/architecturecenter/AWS\\_ac\\_ra\\_disasterrecovery\\_07.pdf](http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_disasterrecovery_07.pdf)

#### NEW QUESTION 187

A government client needs you to set up secure cryptographic key storage for some of their extremely confidential data. You decide that the AWS CloudHSM is the best service for this. However, there seem to be a few pre-requisites before this can happen, one of those being a security group that has certain ports open. Which of the following is correct in regards to those security groups?

- A. A security group that has port 22 (for SSH) or port 3389 (for RDP) open to your network.
- B. A security group that has no ports open to your network.
- C. A security group that has only port 3389 (for RDP) open to your network.
- D. A security group that has only port 22 (for SSH) open to your network.

**Answer:** A

#### Explanation:

AWS CloudHSM provides secure cryptographic key storage to customers by making hardware security modules (HSMs) available in the AWS cloud.

AWS CloudHSM requires the following environment before an HSM appliance can be provisioned. A virtual private cloud (VPC) in the region where you want the AWS CloudHSM service.

One private subnet (a subnet with no Internet gateway) in the VPC. The HSM appliance is provisioned into this subnet.

One public subnet (a subnet with an Internet gateway attached). The control instances are attached to this subnet.

An AWS Identity and Access Management (IAM) role that delegates access to your AWS resources to AWS CloudHSM.

An EC2 instance, in the same VPC as the HSM appliance, that has the SafeNet client software installed. This instance is referred to as the control instance and is used to connect to and manage the HSM appliance.

A security group that has port 22 (for SSH) or port 3389 (for RDP) open to your network. This security group is attached to your control instances so you can access them remotely.

#### NEW QUESTION 188

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

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

**Answer:** D

#### Explanation:

The user has created an Elastic Load Balancer with the availability zone and wants to add more zones to the existing ELB. The user can do so in two ways:

From the console or CLI, add new zones to ELB;

Launch instances in a separate AZ and add instances to the existing ELB. Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/enable-disable-az.html>

#### NEW QUESTION 189

AWS Identity and Access Management is a web service that enables Amazon Web Services (AWS) customers to manage users and user permissions in AWS. In addition to supporting IAM user policies, some services support resource-based permissions. Which of the following services are supported by resource-based permissions?

- A. Amazon SNS, and Amazon SQS and AWS Direct Connect.
- B. Amazon S3 and Amazon SQS and Amazon ElastiCache.
- C. Amazon S3, Amazon SNS, Amazon SQS, Amazon Glacier and Amazon EBS.
- D. Amazon Glacier, Amazon SNS, and Amazon CloudWatch

**Answer:** C

#### Explanation:

In addition to supporting IAM user policies, some services support resource-based permissions, which let you attach policies to the service's resources instead of to IAM users or groups. Resource-based permissions are supported by Amazon S3, Amazon SNS, Amazon SQS, Amazon Glacier and Amazon EBS.

Reference: [http://docs.aws.amazon.com/IAM/latest/UserGuide/Using\\_SpecificProducts.htm](http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_SpecificProducts.htm)

#### NEW QUESTION 194

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

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

**Answer:** A

#### Explanation:

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

Reference: [http://media.amazonwebservices.com/AWS\\_Security\\_Best\\_Practices.pdf](http://media.amazonwebservices.com/AWS_Security_Best_Practices.pdf)

#### NEW QUESTION 195

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

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

**Answer:** C

#### Explanation:

CloudFront is a global service, and metrics are available only when you choose the US East (N. Virginia) region in the AWS console. If you choose another region, no CloudFront metrics will appear in the CloudWatch console.

Reference:

<http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/monitoring-using-cloudwatch.html>

#### NEW QUESTION 200

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

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

**Answer: C**

#### Explanation:

In AWS it is critical to build a scalable architecture in order to take advantage of a scalable infrastructure. The cloud is designed to provide conceptually infinite scalability. However, you cannot leverage all that scalability in infrastructure if your architecture is not scalable. Both have to work together. You will have to identify the monolithic components and bottlenecks in your architecture, identify the areas where you cannot leverage the on-demand provisioning capabilities in your architecture, and work to refactor your application, in order to leverage the scalable infrastructure and take advantage of the cloud.

Characteristics of a truly scalable application:

Increasing resources results in a proportional increase in performance A scalable service is capable of handling heterogeneity

A scalable service is operationally efficient A scalable service is resilient

A scalable service should become more cost effective when it grows (Cost per unit reduces as the number of units increases)

Reference: [http://media.amazonwebservices.com/AWS\\_Cloud\\_Best\\_Practices.pdf](http://media.amazonwebservices.com/AWS_Cloud_Best_Practices.pdf)

#### NEW QUESTION 203

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

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

**Answer: B**

#### Explanation:

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

#### NEW QUESTION 206

Which of the following statements is true of Amazon EC2 security groups?

- A. You can change the outbound rules for EC2-Classic
- B. Also, you can add and remove rules to a group at any time.
- C. You can modify an existing rule in a group
- D. However, you can't add and remove rules to a group.
- E. None of the statements are correct.
- F. You can't change the outbound rules for EC2-Classic
- G. However, you can add and remove rules to a group at any time

**Answer: D**

#### Explanation:

When dealing with security groups, bear in mind that you can freely add and remove rules from a group, but you can't change the outbound rules for EC2-Classic. If you're using the Amazon EC2 console, you can modify existing rules, and you can copy the rules from an existing security group to a new security group.

Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/using-network-security.html>

#### NEW QUESTION 211

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

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

**Answer: D**

#### Explanation:

When the user account has reached the maximum number of EC2 instances, it will not be allowed to launch an instance. AWS will throw an 'Instance Limit Exceeded' error. For all other reasons, such as

"AMI is missing part", "Corrupt Snapshot" or "Volume limit has reached" it will launch an EC2 instance and then terminate it.

Reference: [http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using\\_InstanceStraightToTerminated.html](http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_InstanceStraightToTerminated.html)

#### NEW QUESTION 215

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

- A. Yes, you can change security groups after you launch an instance in EC2-Classic.
- B. No, you cannot change security groups after you launch an instance in EC2-Classic.

- C. Yes, you can only when you remove rules from a security group.
- D. Yes, you can only when you add rules to a security group

**Answer:** B

**Explanation:**

After you launch an instance in EC2-Classic, you can't change its security groups. However, you can add rules to or remove rules from a security group, and those changes are automatically applied to all instances that are associated with the security group.

Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/using-network-security.html>

**NEW QUESTION 217**

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

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

**Answer:** B

**Explanation:**

Each AWS region has multiple, isolated locations known as Availability Zones. To ensure that the AWS resources are distributed across the Availability Zones for a region, AWS independently maps the Availability Zones to identifiers for each account. In this case the Availability Zone US-East-1a where George's EC2 instances are running might not be the same location as the US-East-1a zone of Ray's EC2 instances. There is no way for the user to coordinate the Availability Zones between accounts.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html>

**NEW QUESTION 221**

Can you encrypt EBS volumes?

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

**Answer:** A

**Explanation:**

With Amazon EBS encryption, you can now create an encrypted EBS volume and attach it to a supported instance type. Data on the volume, disk I/O, and snapshots created from the volume are then all encrypted. The encryption occurs on the servers that host the EC2 instances, providing encryption of data as it moves between EC2 instances and EBS storage. EBS encryption is based on the industry standard AES-256 cryptographic algorithm.

To get started, simply enable encryption when you create a new EBS volume using the AWS Management Console, API, or CLI. Amazon EBS encryption is available for all the latest EC2 instances in all commercially available AWS regions.

Reference:

<https://aws.amazon.com/about-aws/whats-new/2014/05/21/Amazon-EBS-encryption-now-available/>

**NEW QUESTION 224**

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

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

**Answer:** A

**Explanation:**

Amazon S3 supports several mechanisms that give you flexibility to control who can access your data as well as how, when, and where they can access it.

Amazon S3 provides four different access control mechanisms:

AWS Identity and Access Management (IAM) policies, Access Control Lists (ACLs), bucket policies, and query string authentication.

IAM enables organizations to create and manage multiple users under a single AWS account. With IAM policies, you can grant IAM users fine-grained control to your Amazon S3 bucket or objects. You can use ACLs to selectively add (grant) certain permissions on individual objects.

Amazon S3 bucket policies can be used to add or deny permissions across some or all of the objects within a single bucket.

With Query string authentication, you have the ability to share Amazon S3 objects through URLs that are valid for a specified period of time.

**NEW QUESTION 226**

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

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

**Answer:** A

**Explanation:**

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

Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/config-conn-drain.html>

**NEW QUESTION 229**

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

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

**Answer:** A

**Explanation:**

In the EC2-Classic platform, when you associate an Elastic IP Address (EIP) with an instance, the instance's current public IP address is released to the EC2-Classic public IP address pool. If you disassociate an EIP from the instance, the instance is automatically assigned a new public IP address within a few minutes. In addition, stopping the instance also disassociates the EIP from it.

But in the EC2-VPC platform, when you associate an EIP with an instance in a default Virtual Private Cloud (VPC), or an instance in which you assigned a public IP to the eth0 network interface during launch, its current public IP address is released to the EC2-VPC public IP address pool. If you disassociate an EIP from the instance, the instance is automatically assigned a new public IP address within a few minutes. However, if you have attached a second network interface to the instance, the instance is not automatically assigned a new public IP address; you'll have to associate an EIP with it manually. The EIP remains associated with the instance when you stop it.

Reference: <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-ip-addresses-eip.html>

**NEW QUESTION 233**

After deciding that EMR will be useful in analysing vast amounts of data for a gaming website that you are architecting you have just deployed an Amazon EMR Cluster and wish to monitor the cluster performance. Which of the following tools cannot be used to monitor the cluster performance?

- A. Kinesis
- B. Ganglia
- C. CloudWatch Metrics
- D. Hadoop Web Interfaces

**Answer:** A

**Explanation:**

Amazon EMR provides several tools to monitor the performance of your cluster. Hadoop Web Interfaces

Every cluster publishes a set of web interfaces on the master node that contain information about the cluster. You can access these web pages by using an SSH tunnel to connect them on the master node. For more information, see [View Web Interfaces Hosted on Amazon EMR Clusters](#).

CloudWatch Metrics

Every cluster reports metrics to CloudWatch. CloudWatch is a web service that tracks metrics, and which you can use to set alarms on those metrics. For more information, see [Monitor Metrics with CloudWatch](#). Ganglia

Ganglia is a cluster monitoring tool. To have this available, you have to install Ganglia on the cluster when you launch it. After you've done so, you can monitor the cluster as it runs by using an SSH tunnel to connect to the Ganglia UI running on the master node. For more information, see [Monitor Performance with Ganglia](#).

Reference:

<http://docs.aws.amazon.com/ElasticMapReduce/latest/DeveloperGuide/emr-troubleshoot-tools.html>

**NEW QUESTION 237**

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

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

**Answer:** C

**Explanation:**

When a user sends a request to ELB over HTTP/HTTPS, the request header log at the instance will only receive the IP of ELB. This is because ELB is the interceptor between the EC2 instance and the client request. To get the client IP, use the header X-Forwarded-For in header. The client IP address in the X-Forwarded-For request header is followed by the IP addresses of each successive proxy that passes along the request. The last IP address is the IP address that connects to the back-end application instance. e.g. if the HTTP request already has a header when it reaches the Load Balancer, the IP address from which the request came is appended at the end of the header followed by the IP address of the Load Balancer. In such cases, the X-Forwarded-For request header takes the following form:

X-Forwarded-For: clientIPAddress, previousRequestIPAddress, LoadBalancerIPAddress. Reference:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/TerminologyandKeyConcepts.html>

**NEW QUESTION 241**

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

- A. createVolumePermission

- B. LaunchPermission
- C. SharePermission
- D. RequestPermission

**Answer:** A

**Explanation:**

Each Amazon EBS Snapshot has a createVolumePermission attribute that you can set to one or more AWS Account IDs to share the AM with those AWS Accounts. To allow several AWS Accounts to use a particular EBS snapshot, you can use the snapshots's createVolumePermission attribute to include a list of the accounts that can use it.

Reference: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/UsingIAM.html>

**NEW QUESTION 244**

A 3-tier e-commerce web application is currently 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 failover 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 time
- B. App servers: share state using a combination of DynamoDB and IP unicast
- 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 time
- F. App servers: share state using a combination of DynamoDB and IP multicast
- 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 time
- J. App servers: share state using a combination of DynamoDB and IP unicast
- 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 time
- N. App servers: share state using a combination of DynamoDB and IP unicast
- O. Database: use RDS with multi-AZ deployment
- P. Backup: web and app servers backed up weekly via AMIs, database backed up via DB snapshots.

**Answer:** C

**Explanation:**

Amazon RDS Multi-AZ deployments provide enhanced availability and durability for Database (DB) Instances, making them a natural fit for production database workloads. When you provision a Multi-AZ DB Instance, Amazon RDS automatically creates a primary DB Instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). Each AZ runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. In case of an infrastructure failure (for example, instance hardware failure, storage failure, or network disruption), Amazon RDS performs an automatic failover to the standby, so that you can resume database operations as soon as the failover is complete. Since the endpoint for your DB Instance remains the same after a failover, your application can resume database operation without the need for manual administrative intervention.

**Benefits**

**Enhanced Durability**

Multi-AZ deployments for the MySQL, Oracle, and PostgreSQL engines utilize synchronous physical replication to keep data on the standby up-to-date with the primary. Multi-AZ deployments for the SQL Server engine use synchronous logical replication to achieve the same result, employing SQL Server-native Mirroring technology. Both approaches safeguard your data in the event of a DB Instance failure or loss of an Availability Zone.

If a storage volume on your primary fails in a Multi-AZ deployment, Amazon RDS automatically initiates a failover to the up-to-date standby. Compare this to a Single-AZ deployment: in case of a Single-AZ database failure, a user-initiated point-in-time-restore operation will be required. This operation can take several hours to complete, and any data updates that occurred after the latest restorable time (typically within the last five minutes) will not be available.

Amazon Aurora employs a highly durable, SSD-backed virtualized storage layer purpose-built for database workloads. Amazon Aurora automatically replicates your volume six ways, across three Availability Zones. Amazon Aurora storage is fault-tolerant, transparently handling the loss of up to two copies of data without affecting database write availability and up to three copies without affecting read availability. Amazon Aurora storage is also self-healing. Data blocks and disks are continuously scanned for errors and replaced automatically.

**Increased Availability**

You also benefit from enhanced database availability when running Multi-AZ deployments. If an Availability Zone failure or DB Instance failure occurs, your availability impact is limited to the time automatic failover takes to complete: typically under one minute for Amazon Aurora and one to two minutes for other database engines (see the RDS FAQ for details).

The availability benefits of Multi-AZ deployments also extend to planned maintenance and backups.

In the case of system upgrades like OS patching or DB Instance scaling, these operations are applied first on the standby, prior to the automatic failover. As a result, your availability impact is, again, only the time required for automatic failover to complete.

Unlike Single-AZ deployments, I/O activity is not suspended on your primary during backup for Multi-AZ deployments for the MySQL, Oracle, and PostgreSQL engines, because the backup is taken from the standby. However, note that you may still experience elevated latencies for a few minutes during backups for Multi-AZ deployments.

On instance failure in Amazon Aurora deployments, Amazon RDS uses RDS Multi-AZ technology to automate failover to one of up to 15 Amazon Aurora Replicas you have created in any of three Availability Zones. If no Amazon Aurora Replicas have been provisioned, in the case of a failure, Amazon RDS will attempt to create a new Amazon Aurora DB instance for you automatically.

**No Administrative Intervention**

DB Instance failover is fully automatic and requires no administrative intervention. Amazon RDS monitors the health of your primary and standbys, and initiates a failover automatically in response to a variety of failure conditions.

**Failover conditions**

Amazon RDS detects and automatically recovers from the most common failure scenarios for Multi-AZ deployments so that you can resume database operations as quickly as possible without administrative intervention. Amazon RDS automatically performs a failover in the event of any of the following:

Loss of availability in primary Availability Zone  
 Loss of network connectivity to primary  
 Compute unit failure on primary

**Storage failure on primary**

Note: When operations such as DB Instance scaling or system upgrades like OS patching are initiated for Multi-AZ deployments, for enhanced availability, they are applied first on the standby prior to an automatic failover. As a result, your availability impact is limited only to the time required for automatic failover to complete.

Note that Amazon RDS Multi-AZ deployments do not failover automatically in response to database operations such as long running queries, deadlocks or database corruption errors.

#### NEW QUESTION 247

A customer has a 10 GB AWS Direct Connect connection to an AWS region where they have a web application hosted on Amazon Elastic Computer Cloud (EC2). The application has dependencies on an on-premises mainframe database that uses a BASE (Basic Available. Sort stale Eventual consistency) rather than an ACID (Atomicity. Consistency isolation. Durability) consistency model. The application is exhibiting undesirable behavior because the database is not able to handle the volume of writes. How can you reduce the load on your on-premises database resources in the most cost-effective way?

- A. Use an Amazon Elastic Map Reduce (EMR) S3DistCp as a synchronization mechanism between the on-premises database and a Hadoop cluster on AWS.
- B. Modify the application to write to an Amazon SQS queue and develop a worker process to flush the queue to the on-premises database.
- C. Modify the application to use DynamoDB to feed an EMR cluster which uses a map function to write to the on-premises database.
- D. Provision an RDS read-replica database on AWS to handle the writes and synchronize the two databases using Data Pipeline.

**Answer:** A

#### Explanation:

Reference: <https://aws.amazon.com/blogs/aws/category/amazon-elastic-map-reduce/>

#### NEW QUESTION 249

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

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

**Answer:** B

#### Explanation:

Amazon Simple Queue Service (Amazon SQS) offers a reliable, highly scalable hosted queue for storing messages as they travel between computers. By using Amazon SQS, developers can simply move data between distributed application components performing different tasks, without losing messages or requiring each component to be always available. Amazon SQS makes it easy to build a distributed, decoupled application, working in close conjunction with the Amazon Elastic Compute Cloud (Amazon EC2) and the other AWS infrastructure web services.

What can I do with Amazon SQS?

Amazon SQS is a web service that gives you access to a message queue that can be used to store messages while waiting for a computer to process them. This allows you to quickly build message queuing applications that can be run on any computer on the internet. Since Amazon SQS is highly scalable and you only pay for what you use, you can start small and grow your application as you wish, with no compromise on performance or reliability. This lets you focus on building sophisticated message-based applications, without worrying about how the messages are stored and managed.

You can use Amazon SQS with software applications in various ways. For example, you can: Integrate Amazon SQS with other AWS infrastructure web services to make applications more reliable and filexible.

Use Amazon SQS to create a queue of work where each message is a task that needs to be completed by a process. One or many computers can read tasks from the queue and perform them. Build a microservices architecture, using queues to connect your microservices.

Keep notifications of significant events in a business process in an Amazon SQS queue. Each event can have a corresponding message in a queue, and applications that need to be aware of the event can read and process the messages.

#### NEW QUESTION 251

You have recently joined a startup company building sensors to measure street noise and air quality in urban areas. The company has been running a pilot deployment of around 100 sensors for 3 months each sensor uploads 1KB of sensor data every minute to a backend hosted on AWS.

During the pilot, you measured a peak of 10 IOPS on the database, and you stored an average of 3GB of sensor data per month in the database.

The current deployment consists of a load-balanced auto scaled Ingestion layer using EC2 instances and a PostgreSQL RDS database with 500GB standard storage.

The pilot is considered a success and your CEO has managed to get the attention of some potential investors. The business plan requires a deployment of at least 1000 sensors which needs to be supported by the backend. You also need to store sensor data for at least two years to be able to compare year over year Improvements.

To secure funding, you have to make sure that the platform meets these requirements and leaves room for further scaling. Which setup will meet the requirements?

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

**Answer:** C

#### NEW QUESTION 252

A web design company currently runs several FTP servers that their 250 customers use to upload and download large graphic files. They wish to move this system to AWS to make it more scalable, but they wish to maintain customer privacy and keep costs to a minimum.

What AWS architecture would you recommend?

- A. Ask their customers to use an S3 client instead of an FTP client
- B. Create a single S3 bucket. Create an IAM user for each customer. Put the IAM Users in a Group that has an IAM policy that permits access to sub-directories within the bucket via use of the 'username' Policy variable.
- C. Create a single S3 bucket with Reduced Redundancy Storage turned on and ask their customers to use an S3 client instead of an FTP client. Create a bucket for each customer with a Bucket Policy that permits access only to that one customer.
- D. Create an auto-scaling group of FTP servers with a scaling policy to automatically scale-in when minimum network traffic on the auto-scaling group is below a given threshold
- E. Load a central list of ftp users from S3 as part of the user Data startup script on each Instance.
- F. Create a single S3 bucket with Requester Pays turned on and ask their customers to use an S3 client instead of an FTP client. Create a bucket for each

customer with a Bucket Policy that permits access only to that one customer.

Answer: A

**NEW QUESTION 254**

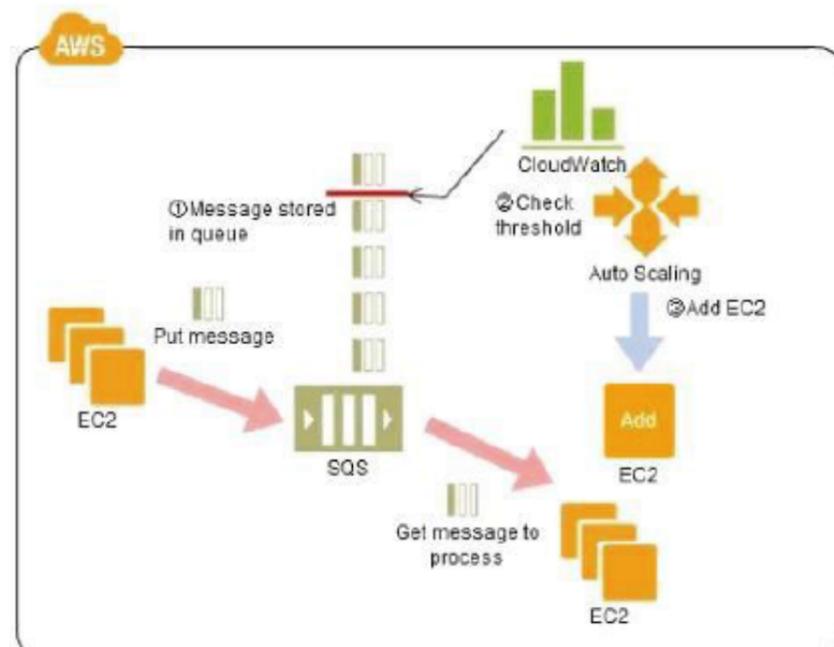
You have been asked to design the storage layer for an application. The application requires disk performance of at least 100,000 IOPS in addition, the storage layer must be able to survive the loss of an individual disk. EC2 instance, or Availability Zone without any data loss. The volume you provide must have a capacity of at least 3 TB. Which of the following designs will meet these objectives'?

- A. Instantiate a c3.8xlarge instance in us-east-1. Provision 4x1TB EBS volumes, attach them to the instance, and configure them as a single RAID 5 volume
- B. Ensure that EBS snapshots are performed every 15 minutes.
- C. Instantiate a c3.8xlarge instance in us-east-1. Provision 3x1TB EBS volumes, attach them to the Instance, and configure them as a single RAID 0 volume
- D. Ensure that EBS snapshots are performed every 15 minutes.
- E. Instantiate an i2.8xlarge instance in us-east-1
- F. Create a RAID 0 volume using the four 800GB SSD ephemeral disks provided with the instance
- G. Provision 3x1TB EBS volumes, attach them to the instance, and configure them as a second RAID 0 volume
- H. Configure synchronous, block-level replication from the ephemeral-backed volume to the EBS-backed volume.
- I. Instantiate a c3.8xlarge instance in us-east-1. Provision an AWS Storage Gateway and configure it for 3 TB of storage and 100,000 IOP
- J. Attach the volume to the instance
- K. Instantiate an i2.8xlarge instance in us-east-1
- L. Create a RAID 0 volume using the four 800GB SSD ephemeral disks provided with the instance
- M. Configure synchronous, block-level replication to an identically configured instance in us-east-1

Answer: C

**NEW QUESTION 257**

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



- A. Reduce the overall time for executing jobs through parallel processing by allowing a busy EC2 instance that receives a message to pass it to the next instance in a daisy-chain setup.
- B. Implement fault tolerance against EC2 instance failure since messages would remain in SQS and work can continue with recovery of EC2 instances implement fault tolerance against SQS failure by backing up messages to S3.
- C. Implement message passing between EC2 instances within a batch by exchanging messages through SQS.
- D. Coordinate number of EC2 instances with number of job requests automatically thus Improving cost effectiveness.
- E. Handle high priority jobs before lower priority jobs by assigning a priority metadata field to SQS messages.

Answer: D

**Explanation:**

Reference:

There are cases where a large number of batch jobs may need processing, and where the jobs may need to be re-prioritized.

For example, one such case is one where there are differences between different levels of services for unpaid users versus subscriber users (such as the time until publication) in services enabling, for example, presentation files to be uploaded for publication from a web browser. When the user uploads a presentation file, the conversion processes, for example, for publication are performed as batch processes on the system side, and the file is published after the conversion. Is it then necessary to be able to assign the level of priority to the batch processes for each type of subscriber.

Explanation of the Cloud Solution/Pattern

A queue is used in controlling batch jobs. The queue need only be provided with priority numbers. Job requests are controlled by the queue, and the job requests in the queue are processed by a batch server. In Cloud computing, a highly reliable queue is provided as a service, which you can use to structure a highly reliable batch system with ease. You may prepare multiple queues depending on priority levels, with job requests put into the queues depending on their priority levels, to apply prioritization to batch processes. The performance (number) of batch servers corresponding to a queue must be in accordance with the priority level thereof.

Implementation

In AWS, the queue service is the Simple Queue Service (SQS). Multiple SQS queues may be prepared to prepare queues for individual priority levels (with a priority queue and a secondary queue).

Moreover, you may also use the message Delayed Send function to delay process execution. Use SQS to prepare multiple queues for the individual priority levels. Place those processes to be executed immediately (job requests) in the high priority queue. Prepare numbers of batch servers, for processing the job requests of

the queues, depending on the priority levels.

Queues have a message "Delayed Send" function. You can use this to delay the time for starting a process.

Configuration

Benefits

You can increase or decrease the number of servers for processing jobs to change automatically the processing speeds of the priority queues and secondary queues.

You can handle performance and service requirements through merely increasing or decreasing the number of EC2 instances used in job processing.

Even if an EC2 were to fail, the messages (jobs) would remain in the queue service, enabling processing to be continued immediately upon recovery of the EC2 instance, producing a system that is robust to failure.

Cautions

Depending on the balance between the number of EC2 instances for performing the processes and the number of messages that are queued, there may be cases where processing in the secondary queue may be completed first, so you need to monitor the processing speeds in the primary queue and the secondary queue.

#### NEW QUESTION 261

You have deployed a web application targeting a global audience across multiple AWS Regions under the domain name.example.com. You decide to use Route53 Latency-Based Routing to serve web requests to users from the region closest to the user. To provide business continuity in the event of server downtime you configure weighted record sets associated with two web servers in separate Availability Zones per region. During a DR test you notice that when you disable all web servers in one of the regions Route53 does not automatically direct all users to the other region. What could be happening? (Choose 2 answers)

- A. Latency resource record sets cannot be used in combination with weighted resource record sets.
- B. You did not setup an HTTP health check for one or more of the weighted resource record sets associated with the disabled web servers.
- C. The value of the weight associated with the latency alias resource record set in the region with the disabled servers is higher than the weight for the other region.
- D. One of the two working web servers in the other region did not pass its HTTP health check.
- E. You did not set "Evaluate Target Health" to "Yes" on the latency alias resource record set associated with example.com in the region where you disabled the servers.

**Answer:** BE

#### Explanation:

How Health Checks Work in Complex Amazon Route 53 Configurations

Checking the health of resources in complex configurations works much the same way as in simple configurations. However, in complex configurations, you use a combination of alias resource record sets (including weighted alias, latency alias, and failover alias) and nonalias resource record sets to build a decision tree that gives you greater control over how Amazon Route 53 responds to requests.

For more information, see How Health Checks Work in Simple Amazon Route 53 Configurations.

For example, you might use latency alias resource record sets to select a region close to a user and use weighted resource record sets for two or more resources within each region to protect against the failure of a single endpoint or an Availability Zone. The following diagram shows this configuration.

Here's how Amazon EC2 and Amazon Route 53 are configured:

You have Amazon EC2 instances in two regions, us-east-1 and ap-southeast-2. You want Amazon Route 53 to respond to queries by using the resource record sets in the region that provides the lowest latency for your customers, so you create a latency alias resource record set for each region.

(You create the latency alias resource record sets after you create resource record sets for the individual Amazon EC2 instances.)

Within each region, you have two Amazon EC2 instances. You create a weighted resource record set for each instance. The name and the type are the same for both of the weighted resource record sets in each region.

When you have multiple resources in a region, you can create weighted or failover resource record sets for your resources. You can also create even more complex configurations by creating weighted alias or failover alias resource record sets that, in turn, refer to multiple resources.

Each weighted resource record set has an associated health check. The IP address for each health check matches the IP address for the corresponding resource record set. This isn't required, but it's the most common configuration.

For both latency alias resource record sets, you set the value of Evaluate Target Health to Yes.

You use the Evaluate Target Health setting for each latency alias resource record set to make Amazon Route 53 evaluate the health of the alias targets-the weighted resource record sets-and respond accordingly.

The preceding diagram illustrates the following sequence of events:

Amazon Route 53 receives a query for example.com. Based on the latency for the user making the request, Amazon Route 53 selects the latency alias resource record set for the us-east-1 region.

Amazon Route 53 selects a weighted resource record set based on weight. Evaluate Target Health is Yes for the latency alias resource record set, so Amazon Route 53 checks the health of the selected weighted resource record set.

The health check failed, so Amazon Route 53 chooses another weighted resource record set based on weight and checks its health. That resource record set also is unhealthy.

Amazon Route 53 backs out of that branch of the tree, looks for the latency alias resource record set with the next-best latency, and chooses the resource record set for ap-southeast-2.

Amazon Route 53 again selects a resource record set based on weight, and then checks the health of the selected resource record set. The health check passed, so Amazon Route 53 returns the applicable value in response to the query.

What Happens When You Associate a Health Check with an Alias Resource Record Set?

You can associate a health check with an alias resource record set instead of or in addition to setting the value of Evaluate Target Health to Yes. However, it's generally more useful if Amazon Route 53 responds to queries based on the health of the underlying resources- the HTTP servers, database servers, and other resources that your alias resource record sets refer to. For example, suppose the following configuration:

You assign a health check to a latency alias resource record set for which the alias target is a group of weighted resource record sets.

You set the value of Evaluate Target Health to Yes for the latency alias resource record set.

In this configuration, both of the following must be true before Amazon Route 53 will return the applicable value for a weighted resource record set:

The health check associated with the latency alias resource record set must pass.

At least one weighted resource record set must be considered healthy, either because it's associated with a health check that passes or because it's not associated with a health check. In the latter case, Amazon Route 53 always considers the weighted resource record set healthy.

If the health check for the latency alias resource record set fails, Amazon Route 53 stops responding to queries using any of the weighted resource record sets in the alias target, even if they're all healthy. Amazon Route 53 doesn't know the status of the weighted resource record sets because it never looks past the failed health check on the alias resource record set.

What Happens When You Omit Health Checks?

In a complex configuration, it's important to associate health checks with all of the non-alias resource record sets. Let's return to the preceding example, but assume that a health check is missing on one of the weighted resource record sets in the us-east-1 region:

Here's what happens when you omit a health check on a non-alias resource record set in this configuration:

Amazon Route 53 receives a query for example.com. Based on the latency for the user making the request, Amazon Route 53 selects the latency alias resource record set for the us-east-1 region.

Amazon Route 53 looks up the alias target for the latency alias resource record set, and checks the status of the corresponding health checks. The health check for one weighted resource record set failed, so that resource record set is omitted from consideration.

The other weighted resource record set in the alias target for the us-east-1 region has no health check. The corresponding resource might or might not be healthy, but without a health check, Amazon Route 53 has no way to know. Amazon Route 53 assumes that the resource is healthy and returns the applicable value in response to the query.

What Happens When You Set Evaluate Target Health to No?

In general, you also want to set Evaluate Target Health to Yes for all of the alias resource record sets. In the following example, all of the weighted resource record sets have associated health checks, but Evaluate Target Health is set to No for the latency alias resource record set for the us-east-1 region:

Here's what happens when you set Evaluate Target Health to No for an alias resource record set in this configuration:

Amazon Route 53 receives a query for example.com. Based on the latency for the user making the request, Amazon Route 53 selects the latency alias resource record set for the us-east-1 region.

Amazon Route 53 determines what the alias target is for the latency alias resource record set, and checks the corresponding health checks. They're both failing. Because the value of Evaluate Target Health is No for the latency alias resource record set for the us-east-1 region, Amazon Route 53 must choose one resource record set in this branch instead of backing out of the branch and looking for a healthy resource record set in the ap-southeast-2 region.

#### NEW QUESTION 263

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: D

#### NEW QUESTION 268

.....

## THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AWS-Solution-Architect-Associate Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AWS-Solution-Architect-Associate Product From:

<https://www.2passeasy.com/dumps/AWS-Solution-Architect-Associate/>

## Money Back Guarantee

### **AWS-Solution-Architect-Associate Practice Exam Features:**

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