

## SAA-C01 Dumps

### AWS Certified Solutions Architect - Associate

<https://www.certleader.com/SAA-C01-dumps.html>



**NEW QUESTION 1**

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 throughput.

**Answer: B**

**Explanation:**

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

There is no limit on the number of messages that can be pushed onto SQS. The retention period of the SQS is 4 days by default and it can be changed to 14 days. This will make sure that no writes are missed.

**NEW QUESTION 2**

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 100K 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 10K provisioned IOPS

**Answer: C**

**Explanation:**

You cannot go with DynamoDB because the application is currently using a PostgreSQL which is an RDS. Replacing an RDS SQL with a noSQL DB, for the sake of scaling is not a sensible option.

**NEW QUESTION 3**

Your application is using an ELB in front of an Auto Scaling group of web/application servers deployed across two AZs and a Multi-AZ RDS Instance for data persistence.

The database CPU is often above 80% usage and 90% of I/O operations on the database are reads. To improve performance you recently added a single-node Memcached ElastiCache Cluster to cache frequent DB query results. In the next weeks the overall workload is expected to grow by 30%.

Do you need to change anything in the architecture to maintain the high availability of the application with the anticipated additional load? Why?

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

**Answer: A**

**Explanation:**

A single-node Memcached ElastiCache cluster failure is nothing but a total failure. (Even though AWS will automatically recover the failed node, there are no other nodes in the cluster) <http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/BestPractices.html> Mitigating Node Failures

To mitigate the impact of a node failure, spread your cached data over more nodes. Because Memcached does not support replication, a node failure will always result in some data loss from your cluster.

When you create your Memcached cluster you can create it with 1 to 20 nodes, or more by special request. Partitioning your data across a greater number of nodes means you'll lose less data if a node fails. For example, if you partition your data across 10 nodes, any single node stores approximately 10% of your cached data. In this case, a node failure loses approximately 10% of your cache which needs to be replaced when a replacement node is created and provisioned. Mitigating Availability Zone Failures

To mitigate the impact of an availability zone failure, locate your nodes in as many availability zones as possible. In the unlikely event of an AZ failure, you will lose only the data cached in that AZ, not the data cached in the other AZs.

**NEW QUESTION 4**

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

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

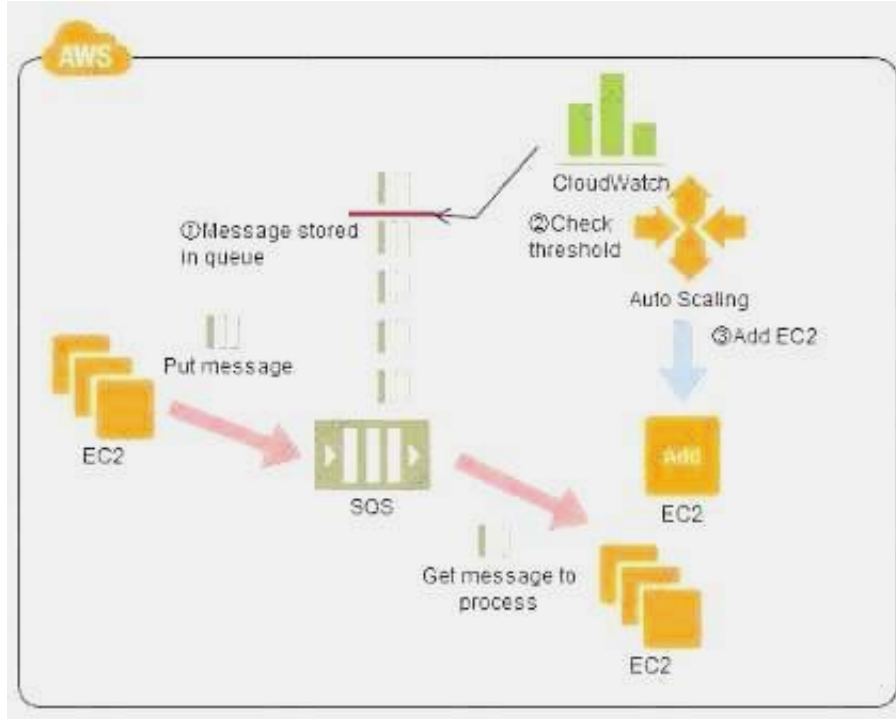
- A. Use AWS Data Pipeline to schedule a DynamoDB cross region copy once a day
- B. Create a 'LastUpdated' attribute in your DynamoDB table that would represent the timestamp of the last update and use it as a filter.
- C. Use EMR and write a custom script to retrieve data from DynamoDB in the current region using a SCAN operation and push it to DynamoDB in the second region.

- D. Use AWS data Pipeline to schedule an export of the DynamoDB table to S3 in the current region once a day then schedule another task immediately after it that will import data from S3 to DynamoDB in the other region.
- E. Send also each Ante into an SQS queue in me second region; use an auto-scaling group behind the SQS queue to replay the write in the second region.

**Answer:** A

#### NEW QUESTION 5

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 lime for executing jobs through parallel processing by allowing a busy EC2 instance that receives a message to pass it to the next instance in a daisy-chain setup.
- B. Implement fault tolerance against EC2 instance failure since messages would remain in SQS and worn can continue with recovery of EC2 instances implement fault tolerance against SQS failure by backing up messages to 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

#### NEW QUESTION 6

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

- A. Take hourly DB backups to S3, with transaction logs stored in S3 every 5 minutes.
- B. Use synchronous database master-slave replication between two availability zones.
- C. Take hourly DB backups to EC2 Instance store volumes with transaction logs stored In S3 every 5 minutes.
- D. Take 15-minute DB backups stored In Glacier with transaction logs stored in S3 every 5 minute

**Answer:** A

#### NEW QUESTION 7

Your startup wants to implement an order fulfillment process for selling a personalized gadget that needs an average of 3-4 days to produce with some orders taking up to 6 months you expect 10 orders per day on your first day. 1000 orders per day after 6 months and 10,000 orders after 12 months.

Orders coming in are checked for consistency men dispatched to your manufacturing plant for production quality control packaging shipment and payment processing If the product does not meet the quality standards at any stage of the process employees may force the process to repeat a step Customers are notified via email about order status and any critical issues with their orders such as payment failure.

Your case architecture includes AWS Elastic Beanstalk for your website with an RDS MySQL instance for customer data and orders.

How can you implement the order fulfillment process while making sure that the emails are delivered reliably?

- A. Add a business process management application to your Elastic Beanstalk app servers and re-use the ROS database for tracking order status use one of the Elastic Beanstalk instances to send emails to customers.
- B. Use SWF with an Auto Scaling group of activity workers and a decider instance in another Auto Scaling group with min/max=1 Use the decider instance to send emails to customers.
- C. Use SWF with an Auto Scaling group of activity workers and a decider instance in another Auto Scaling group with min/max=1 use SES to send emails to customers.
- D. Use an SQS queue to manage all process tasks Use an Auto Scaling group of EC2 Instances that poll the tasks and execute the
- E. Use SES to send emails to customers.

**Answer:** C

#### Explanation:

[http://media.amazonwebservices.com/architecturecenter/AWS\\_ac\\_ra\\_ecommerce\\_checkout\\_13.pdf](http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_ecommerce_checkout_13.pdf)

## NEW QUESTION 8

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. Running 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 two.)

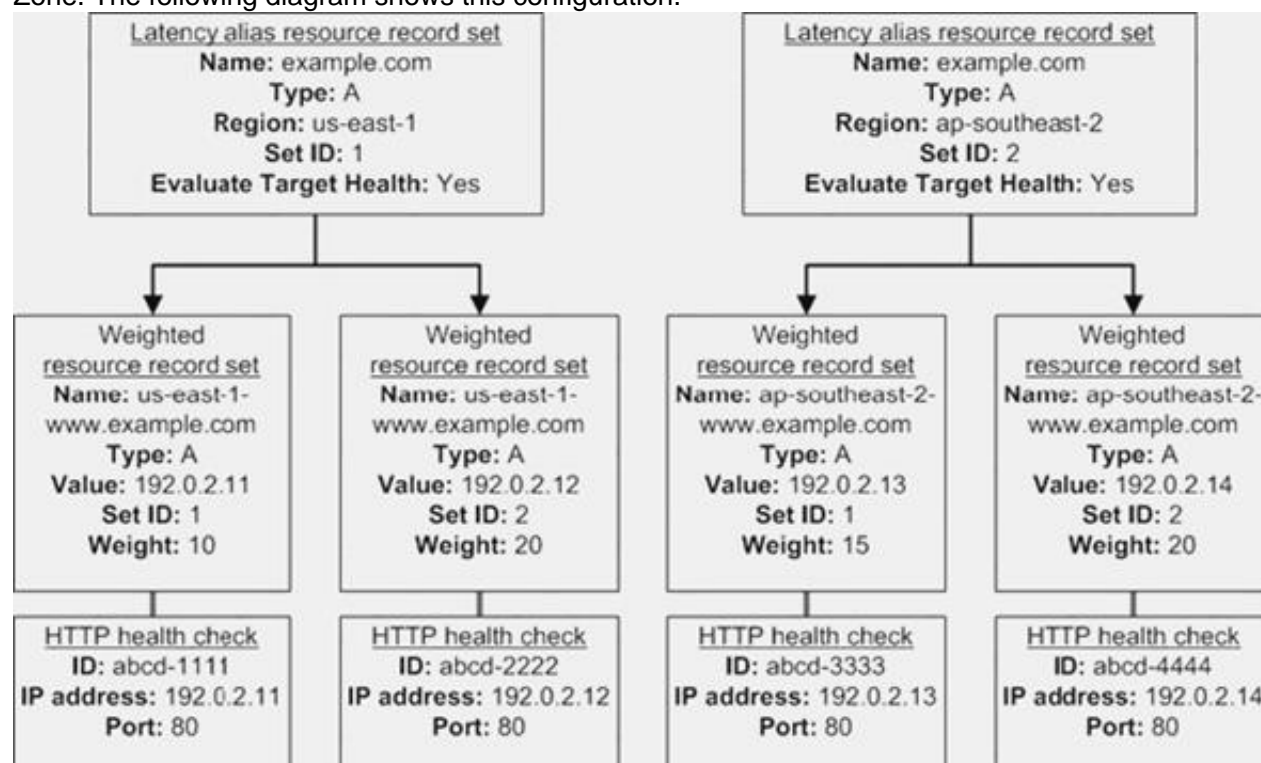
- A. Latency resource record sets cannot be used in combination with weighted resource record sets.
- B. You did not setup an HTTP health check for one or more of the weighted resource record sets associated with the disabled web servers.
- C. The value of the weight associated with the latency alias 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.

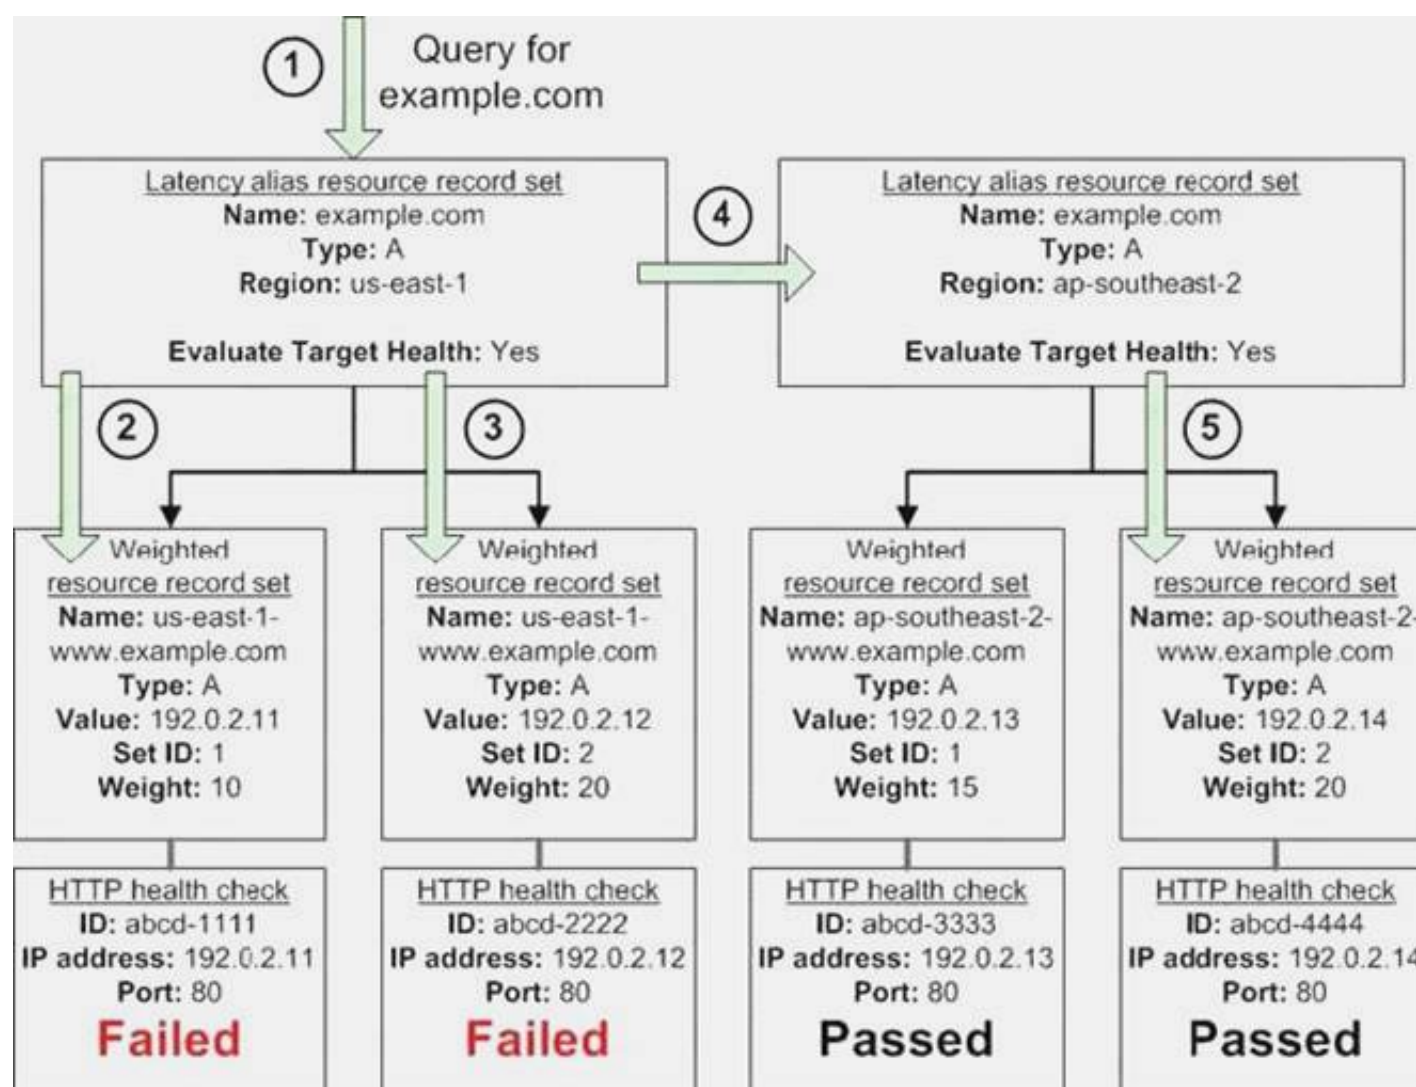
When you have multiple resources in a region, you can create weighted or failover resource record sets for your resources. You can also create even more complex configurations by creating weighted alias or failover alias resource record sets that, in turn, refer to multiple resources.

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:

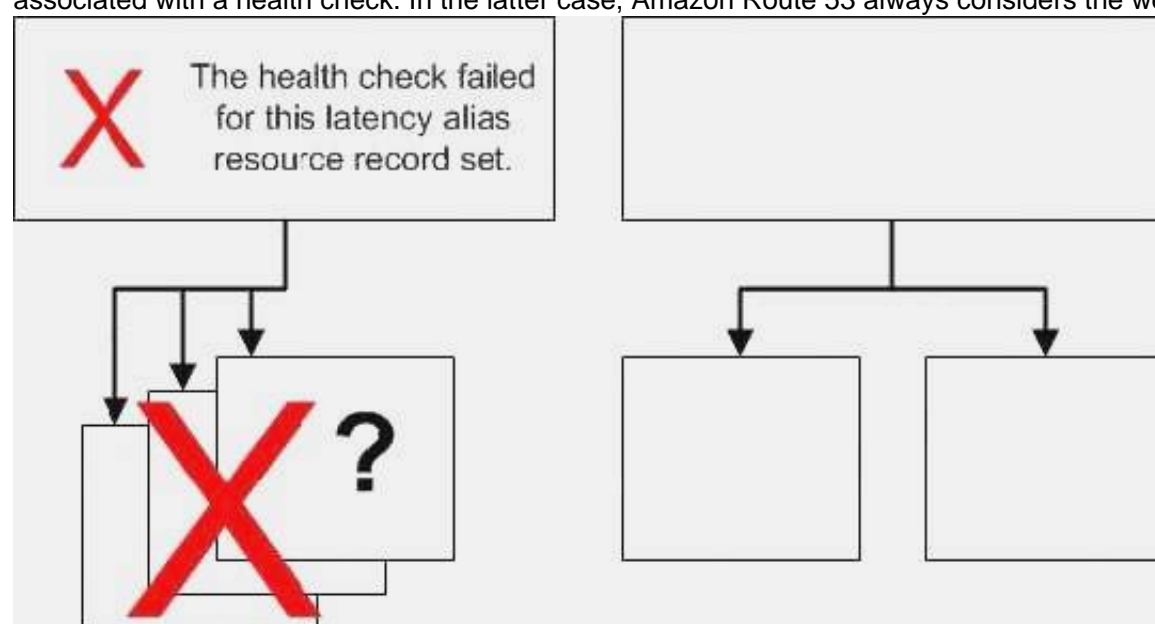
You assign a health check to a latency alias resource record set for which the alias target is a group of weighted resource record sets.

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.

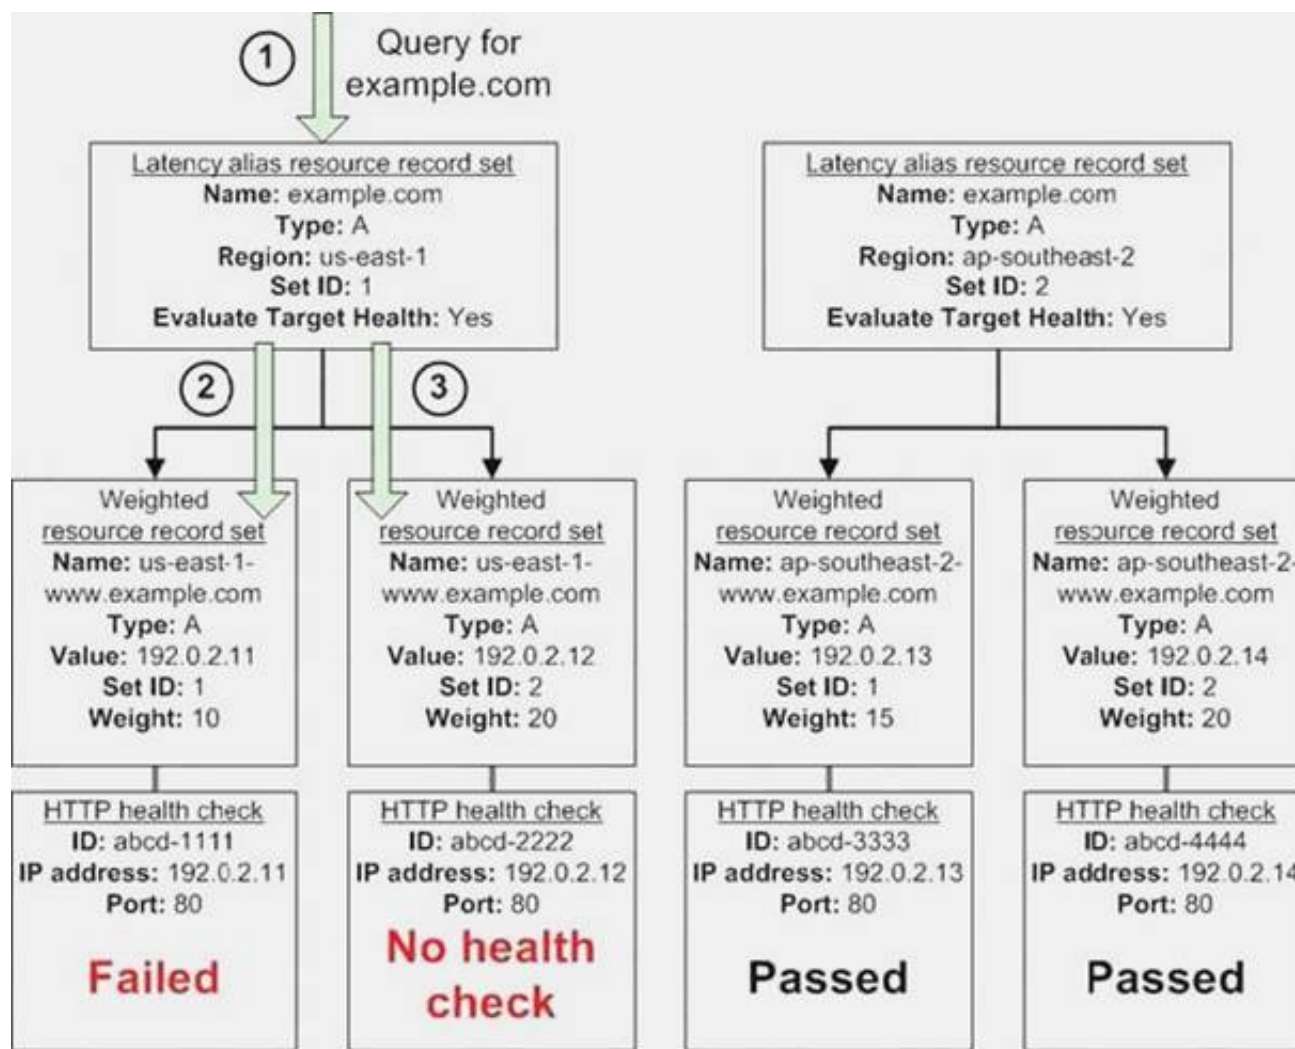
At least one weighted resource record set must be considered healthy, either because it's associated with a health check that passes or because it's not associated with a health check. In the latter case, Amazon Route 53 always considers the weighted resource record set healthy.



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:



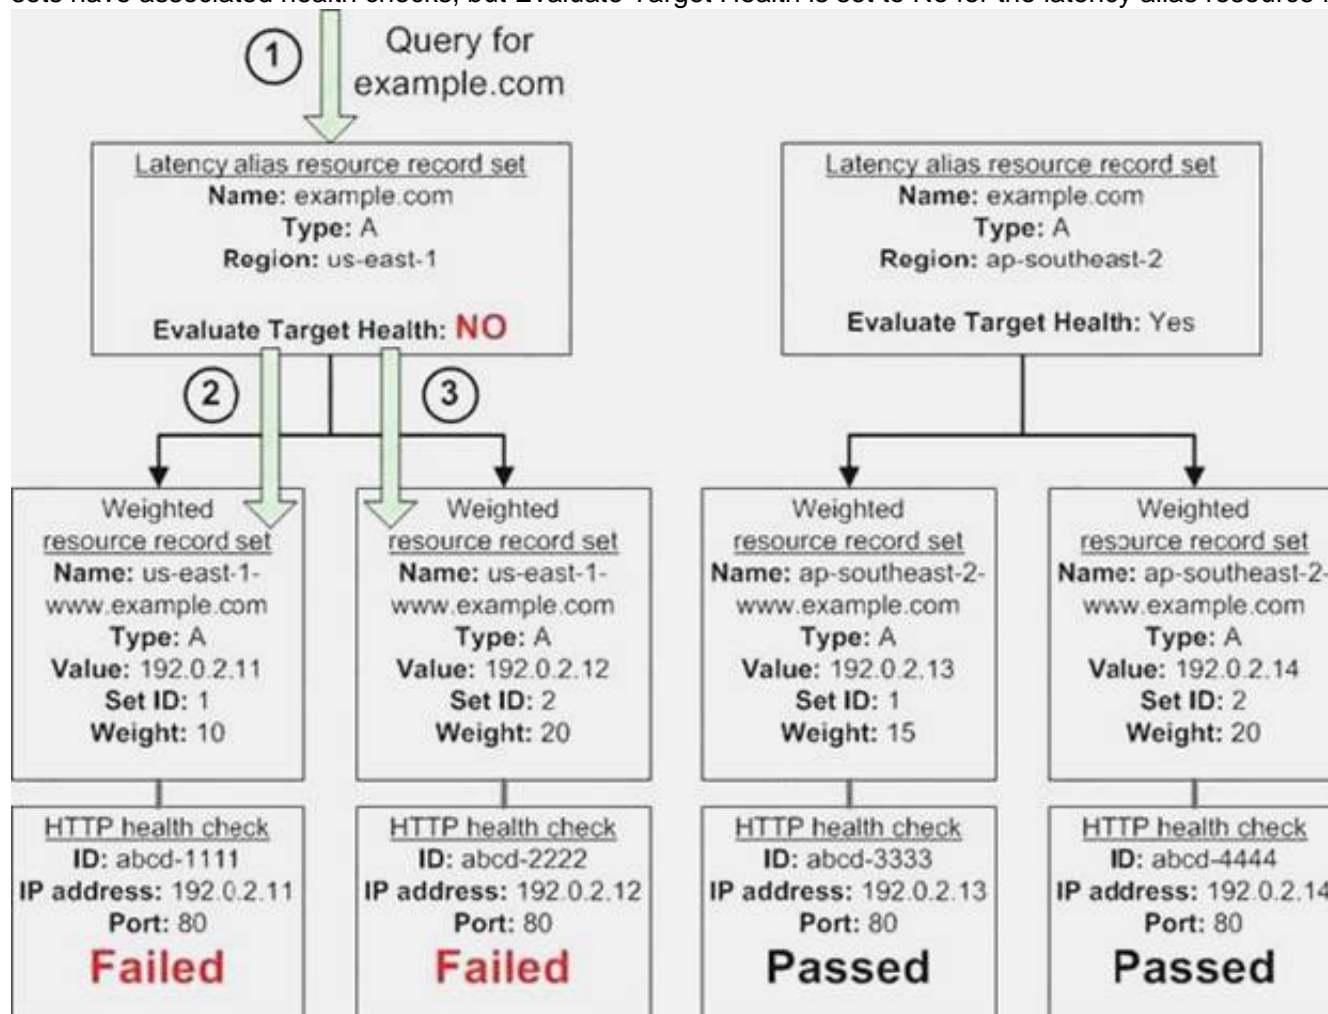
Here's what happens when you omit a health check on a non-alias resource record set in this configuration:

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 useast- 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 9

A web company is looking to implement an external payment service into their highly available application deployed in a VPC. Their application EC2 instances are behind a public facing ELB Auto scaling is used to add additional instances as traffic increases under normal load the application runs 2 instances in the Auto Scaling group but at peak it can scale 3x in size. The application instances need to communicate with the payment service over the Internet which requires



whitelisting of all public IP addresses used to communicate with it. A maximum of 4 whitelisting IP addresses is allowed at a time and can be added through an API.

How should they architect their solution?

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

**Answer:** A

**Explanation:**

B is incorrect as you do not have insight into the public IP associated with a VPC Internet Gateway. C is incorrect as ELB receives a public DNS name. D would exceed the maximum of 4 whitelisting IP addresses.

**NEW QUESTION 10**

You are designing the network infrastructure for an application server in Amazon VPC. Users will access all the application instances from the Internet as well as from an on-premises network. The on-premises network is connected to your VPC over an AWS Direct Connect link.

How would you design routing to meet the above requirements?

- A. Configure a single routing table with a default route via the Internet gateway. Propagate a default route via BGP on the AWS Direct Connect customer route.
- B. Associate the routing table with all VPC subnets.
- C. Configure a single routing table with a default route via the Internet gateway. Propagate specific routes for the on-premises networks via BGP on the AWS Direct Connect customer router. Associate the routing table with all VPC subnets.
- D. Configure a single routing table with two default routes: one to the Internet via an Internet gateway, the other to the on-premises network via the VPN gateway. Use this routing table across all subnets in your VPC.
- E. Configure two routing tables: one that has a default route via the Internet gateway and another that has a default route via the VPN gateway. Associate both routing tables with each VPC subnet.

**Answer:** B

**NEW QUESTION 10**

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

- A. Use S3 with reduced redundancy to store and serve the scanned files, install the commercial search application on EC2 instances and configure with auto-scaling and an Elastic Load Balancer.
- B. Model the environment using CloudFormation, use an EC2 instance running Apache webserver and an open source search application, stripe multiple standard EBS volumes together to store the JPEGs and search index.
- C. Use S3 with standard redundancy to store and serve the scanned files, use CloudSearch for query processing, and use Elastic Beanstalk to host the website across multiple availability zones.
- D. Use a single-AZ RDS MySQL instance to store the search index. Store the JPEG images on EC2 instances to serve the website and translate user queries into SQL.
- E. Use a CloudFront download distribution to serve the JPEGs to the end users and install the current commercial search product, along with a Java container. Run the website on EC2 instances and use Route53 with DNS round-robin.

**Answer:** C

**Explanation:**

There is no such thing as "Most appropriate" without knowing all your goals. I find your scenarios very fuzzy, since you can obviously mix-n-match between them. I think you should decide by layers instead:

Load Balancer Layer: ELB or just DNS, or roll-your-own. (Using DNS+EIPs is slightly cheaper, but less reliable than ELB.)

Storage Layer for 17TB of Images: This is the perfect use case for S3. Off-load all the web requests directly to the relevant JPEGs in S3. Your EC2 boxes just generate links to them.

If your app already serves its own images (not links to images), you might start with EFS. But more than likely, you can just setup a web server to re-write or re-direct all JPEG links to S3 pretty easily. If you use S3, don't serve directly from the bucket - Serve via a CNAME in domain you control. That way, you can switch in CloudFront easily.

EBS will be way more expensive, and you'll need 2x the drives if you need 2 boxes. Yuck. Consider a smaller storage format. For example, JPEG2000 or WebP or other tools might make for smaller images. There is also the DejaVu format from a while back.

Cache Layer: Adding CloudFront in front of S3 will help people on the other side of the world -- well, possibly. Typical archives follow a power law. The long tail of requests means that most JPEGs won't be requested enough to be in the cache. So you are only speeding up the most popular objects. You can always wait, and switch in CF later after you know your costs better. (In some cases, it can actually lower costs.)

You can also put CloudFront in front of your app, since your archive search results should be fairly static. This will also allow you to run with a smaller instance type, since CF will handle much of the load if you do it right.

Database Layer: A few options:

Use whatever your current server does for now, and replace with something else down the road. Don't under-estimate this approach, sometimes it's better to start now and optimize later.

Use RDS to run MySQL/Postgres

I'm not as familiar with ElasticSearch / Cloudsearch, but obviously Cloudsearch will be less maintenance+setup.

App Layer:

When creating the app layer from scratch, consider CloudFormation and/or OpsWorks. It's extra stuff to learn, but helps down the road.

Java+Tomcat is right up the alley of ElasticBeanstalk. (Basically EC2 + Autoscale + ELB). Preventing Abuse: When you put something in a public S3 bucket, people will hot-link it from their web pages. If you want to prevent that, your app on the EC2 box can generate signed links to S3 that expire in a few hours. Now everyone will be forced to go thru the app, and the app can apply rate limiting, etc.

Saving money: If you don't mind having downtime:

run everything in one AZ (both DBs and EC2s). You can always add servers and AZs down the road, as long as it's architected to be stateless. In fact, you should use multiple regions if you want it to be

really robust.

use Reduced Redundancy in S3 to save a few hundred bucks per month (Someone will have to "go fix it" every time it breaks, including having an off-line copy to repair S3.)

Buy Reserved Instances on your EC2 boxes to make them cheaper. (Start with the RI market and buy a partially used one to get started.) It's just a coupon saying "if you run this type of box in this AZ, you will save on the per-hour costs." You can get 1/2 to 1/3 off easily.

Rewrite the application to use less memory and CPU - that way you can run on fewer/smaller boxes. (May or may not be worth the investment.)

If your app will be used very infrequently, you will save a lot of money by using Lambda. I'd be worried that it would be quite slow if you tried to run a Java application on it though.

We're missing some information like load, latency expectations from search, indexing speed, size of the search index, etc. But with what you've given us, I would go with S3 as the storage for the files (S3 rocks. It is really, really awesome). If you're stuck with the commercial search application, then on EC2 instances with autoscaling and an ELB. If you are allowed an alternative search engine, Elasticsearch is probably your best bet. I'd run it on EC2 instead of the AWS Elasticsearch service, as IMHO it's not ready yet. Don't autoscale Elasticsearch automatically though, it'll cause all sorts of issues. I have zero experience with CloudSearch so I can't comment on that. Regardless of which option, I'd use CloudFormation for all of it.

#### NEW QUESTION 15

You've been hired to enhance the overall security posture for a very large e-commerce site. They have a well architected multi-tier application running in a VPC that uses ELBs in front of both the web and the app tier with static assets served directly from S3. They are using a combination of RDS and DynamoDB for their dynamic data and then archiving nightly into S3 for further processing with EMR. They are concerned because they found questionable log entries and suspect someone is attempting to gain unauthorized access.

Which approach provides a cost effective scalable mitigation to this kind of attack?

- A. Recommend that they lease space at a DirectConnect partner location and establish a 1G DirectConnect connection to their VPC. They would then establish Internet connectivity into their space, filter the traffic in hardware Web Application Firewall (WAF). And then pass the traffic through the DirectConnect connection into their application running in their VPC.
- B. Add previously identified hostile source IPs as an explicit INBOUND DENY NACL to the web tier subnet.
- C. Add a WAF tier by creating a new ELB and an AutoScaling group of EC2 Instances running a host-based WAF. They would redirect Route 53 to resolve to the new WAF tier ELB. The WAF tier would then pass the traffic to the current web tier. The web tier Security Groups would be updated to only allow traffic from the WAF tier Security Group.
- D. Remove all but TLS 1.2 from the web tier ELB and enable Advanced Protocol Filtering. This will enable the ELB itself to perform WAF functionality.

**Answer: C**

#### NEW QUESTION 18

An AWS customer is deploying an application that is composed of an AutoScaling group of EC2 Instances.

The customer's security policy requires that every outbound connection from these instances to any other service within the customer's Virtual Private Cloud must be authenticated using a unique x.509 certificate that contains the specific instance-id.

In addition, an x.509 certificate must be designed by the customer's Key management service in order to be trusted for authentication.

Which of the following configurations will support these requirements?

- A. Configure an IAM Role that grants access to an Amazon S3 object containing a signed certificate and configure the Auto Scaling group to launch instances with this role. Have the instances bootstrap get the certificate from Amazon S3 upon first boot.
- B. Embed a certificate into the Amazon Machine Image that is used by the Auto Scaling group. Have the launched instances generate a certificate signature request with the instance's assigned instance-id to the Key management service for signature.
- C. Configure the Auto Scaling group to send an SNS notification of the launch of a new instance to the trusted key management service.
- D. Have the Key management service generate a signed certificate and send it directly to the newly launched instance.
- E. Configure the launched instances to generate a new certificate upon first boot. Have the Key management service poll the AutoScaling group for associated instances and send new instances a certificate signature (that contains the specific instance-id).

**Answer: A**

#### Explanation:

<http://jayendrapatil.com/tag/iam/>

#### NEW QUESTION 19

You are implementing a URL whitelisting system for a company that wants to restrict outbound HTTP/S connections to specific domains from their EC2-hosted applications. You deploy a single EC2 instance running proxy software and configure it to accept traffic from all subnets and EC2 instances in the VPC. You configure the proxy to only pass through traffic to domains that you define in its whitelist configuration. You have a nightly maintenance window of 10 minutes where all instances fetch new software updates. Each update is about 200MB in size and there are 500 instances in the VPC that routinely fetch updates. After a few days you notice that some machines are failing to successfully download some, but not all of their updates within the maintenance window. The download URLs used for these updates are correctly listed in the proxy's whitelist configuration and you are able to access them manually using a web browser on the instances. What might be happening? (Choose two.)

- A. You are running the proxy on an undersized EC2 instance type so network throughput is not sufficient for all instances to download their updates in time.
- B. You are running the proxy on a sufficiently-sized EC2 instance in a private subnet and its network throughput is being throttled by a NAT running on an undersized EC2 instance.
- C. The route table for the subnets containing the affected EC2 instances is not configured to direct network traffic for the software update locations to the proxy.
- D. You have not allocated enough storage to the EC2 instance running the proxy so the network buffer is filling up, causing some requests to fail.
- E. You are running the proxy in a public subnet but have not allocated enough EIPs to support the needed network throughput through the Internet Gateway (IGW).

**Answer: AB**

#### NEW QUESTION 21

To serve Web traffic for a popular product your chief financial officer and IT director have purchased 10 m1 large heavy utilization Reserved Instances (RIs) evenly spread across two availability zones. Route 53 is used to deliver the traffic to an Elastic Load Balancer (ELB). After several months, the product grows even more popular and you need additional capacity. As a result, your company purchases two C3.2xlarge medium utilization RIs.

You register the two c3.2xlarge instances with your ELB and quickly find that the m1 large instances are at 100% of capacity and the c3.2xlarge instances have significant capacity that's unused.



Which option is the most cost effective and uses EC2 capacity most effectively?

- A. Use a separate ELB for each instance type and distribute load to ELBs with Route 53 weighted round robin
- B. Configure Autoscaling group and Launch Configuration with ELB to add up to 10 more on-demand m1 large instances when triggered by Cloudwatch shut off c3 2xlarge instances
- C. Route traffic to EC2 m1 large and c3 2xlarge instances directly using Route 53 latency based routing and health checks shut off ELB
- D. Configure ELB with two c3 2xlarge Instances and use on-demand Autoscaling group for up to two additional c3.2xlarge instances Shut on m1 .large instances.

**Answer:** A

**Explanation:**

Weighted Routing Policy

Use the weighted routing policy when you have multiple resources that perform the same function (for example, web servers that serve the same website) and you want Amazon Route 53 to route traffic to those resources in proportions that you specify (for example, one quarter to one server and three quarters to the other). For more information about weighted resource record sets, see Weighted Routing.

**NEW QUESTION 24**

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

What is the best approach to meet your customer's requirements?

- A. Send all the log events to Amazon SQ
- B. Setup an Auto Scaling group of EC2 servers to consume the logs and apply the heuristics.
- C. Send all the log events to Amazon Kinesis develop a client process to apply heuristics on the logs
- D. Configure Amazon Cloud Trail to receive custom logs, use EMR to apply heuristics the logs
- E. Setup an Auto Scaling group of EC2 syslogd servers, store the logs on S3 use EMR to apply heuristics on the logs

**Answer:** B

**Explanation:**

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

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

**NEW QUESTION 26**

What does Amazon SWF stand for?

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

**Answer:** B

**NEW QUESTION 28**

If I write the below command, what does it do? `ec2-run ami-e3a5408a -n 20 -g appserver`

- A. Start twenty instances as members of appserver group.
- B. Creates 20 rules in the security group named appserver
- C. Terminate twenty instances as members of appserver group.
- D. Start 20 security groups

**Answer:** A

**NEW QUESTION 29**

Every user you create in the IAM system starts with .

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

**Answer:** C

**NEW QUESTION 30**

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

- A. No
- B. Only in VPC
- C. Yes

**Answer:** C

**NEW QUESTION 34**

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

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

**Answer:** C

**NEW QUESTION 38**

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

- A. FreeStorage
- B. FreeStorageSpace
- C. FreeStorageVolume
- D. FreeDBStorageSpace

**Answer:** B

**NEW QUESTION 40**

Groups can't \_\_\_\_.

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 42**

Which is the default region in AWS?

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 47**

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

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

**Answer:** B

**Explanation:**

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

You can run Amazon RDS for Oracle under two different licensing models – **"License Included"** and **"Bring-Your-Own-License (BYOL)"**. In the "License Included" service model, you do not need separately purchased Oracle licenses; the Oracle Database software has been licensed by AWS. "License Included" pricing starts at \$0.04 per hour, inclusive of software, underlying hardware resources, and Amazon RDS management capabilities. If you already own Oracle Database licenses, you can use the "BYOL" model to run Oracle databases on Amazon RDS, with rates starting at \$0.025 per hour. The "BYOL" model is designed for customers who prefer to use existing Oracle database licenses or purchase new licenses directly from Oracle. For more information, see [Licensing Amazon RDS for Oracle](#).

**NEW QUESTION 50**

Which of the following cannot be used in Amazon EC2 to control who has access to specific Amazon EC2 instances?

- A. Security Groups
- B. IAM System
- C. SSH keys
- D. Windows passwords

**Answer:** B

**Explanation:**

<http://blogs.aws.amazon.com/security/post/Tx29HCT3ABL7LP3/Resource-level-Permissions-for-EC2-Controlling-Management-Access-on-Specific-Ins>

**NEW QUESTION 52**

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

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

**Answer:** B

**NEW QUESTION 54**

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

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

**Answer:** A

**Explanation:**

Currently you can request to increase the limit on users per AWS account, groups per AWS account, roles per AWS account, instance profiles per AWS account, and server certificates per AWS account.

This never states “groups a user can be in”

**NEW QUESTION 58**

Will my standby RDS instance be in the same Availability Zone as my primary?

- A. Only for Oracle RDS types
- B. Yes
- C. Only if configured at launch
- D. No

**Answer:** D

**NEW QUESTION 61**

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

- A. Yes but only in certain cases
- B. No
- C. Yes always

**Answer:** C

**Explanation:**

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

**NEW QUESTION 63**

Are Reserved Instances available for Multi-AZ Deployments?

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 64**

If I modify a DB Instance or the DB parameter group associated with the instance, should I reboot the instance for the changes to take effect?

- A. No
- B. Yes

**Answer:** B

**NEW QUESTION 69**

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

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



**Answer:** A

**Explanation:**

If a source DB instance has several Read Replicas, promoting one of the Read Replicas to a DB instance has no effect on the other replicas.

**NEW QUESTION 72**

What does Amazon CloudFormation provide?

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

**Answer:** C

**Explanation:**

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/Welcome.html>

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.

**NEW QUESTION 73**

Can the string value of 'Key' be prefixed with: aws:"?"

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

**Answer:** D

**Explanation:**

"The tag key is the required name of the tag. The string value can be from 1 to 128 Unicode characters in length and cannot be prefixed with "aws:" or "rds:." "

[http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER\\_Tagging.html](http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_Tagging.html) <http://docs.aws.amazon.com/cli/latest/reference/rds/list-tags-for-resource.html>

**NEW QUESTION 74**

How are the EBS snapshots saved on Amazon S3?

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

**Answer:** B

**NEW QUESTION 76**

What happens when you create a topic on Amazon SNS?

- A. The topic is created, and it has the name you specified for it.
- B. An ARN (Amazon Resource Name) is created.
- C. You can create a topic on Amazon SQS, not on Amazon SNS.
- D. This question doesn't make sens

**Answer:** B

**NEW QUESTION 78**

Do the Amazon EBS volumes persist independently from the running life of an Amazon EC2 instance?

- A. Only if instructed to when created
- B. Yes
- C. No

**Answer:** B

**Explanation:**

Data persistence

An EBS volume is off-instance storage that can persist independently from the life of an instance. You continue to pay for the volume usage as long as the data persists.

References:

**NEW QUESTION 80**

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

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

D. Role

**Answer:** A

**Explanation:**

Q: What is a user?

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

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

**NEW QUESTION 81**

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

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

**Answer:** C

**Explanation:**

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

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

**NEW QUESTION 82**

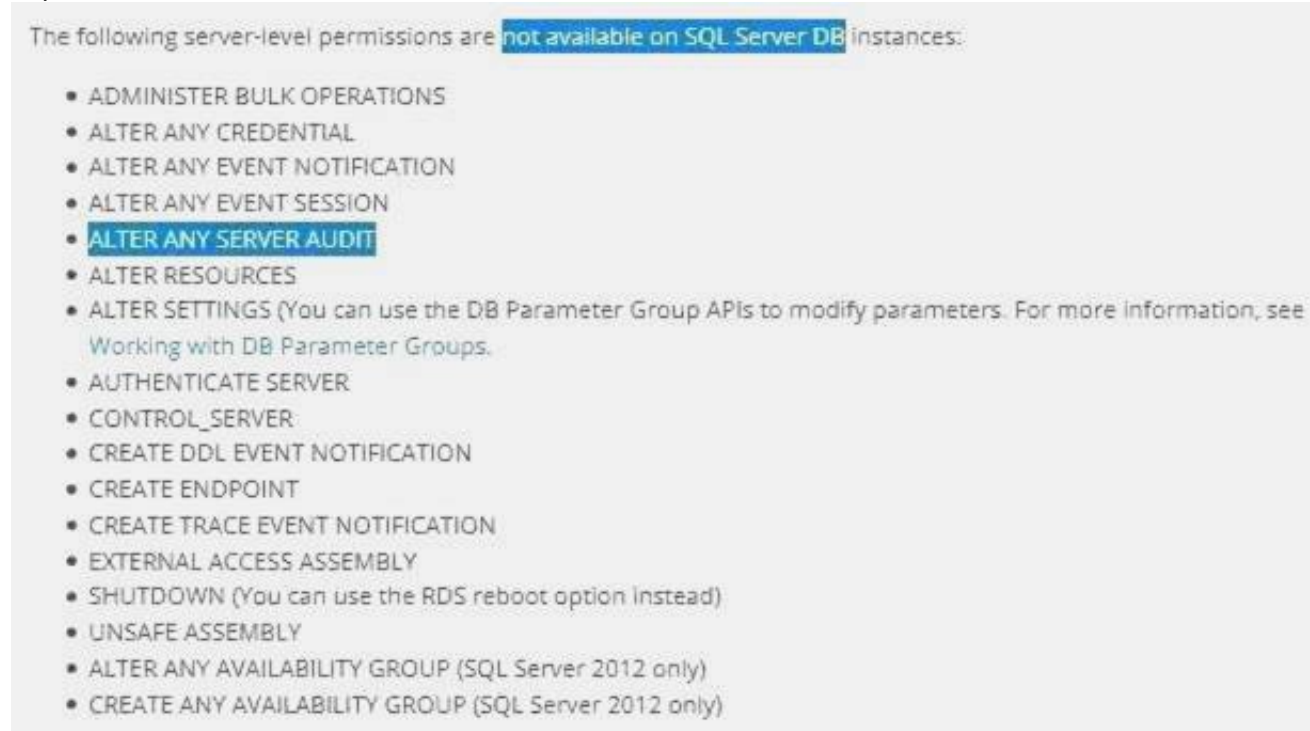
Is the SQL Server Audit feature supported in the Amazon RDS SQL Server engine?

- A. No
- B. Yes

**Answer:** A

**Explanation:**

[http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP\\_SQLServer.html](http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_SQLServer.html)



**NEW QUESTION 87**

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

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

**Answer:** A

**NEW QUESTION 90**

Do the Amazon EBS volumes persist independently from the running life of an Amazon EC2 instance?

- A. No
- B. Only if instructed to when created
- C. Yes

**Answer:** C

**NEW QUESTION 94**

If I want my instance to run on a single-tenant hardware, which value do I have to set the instance's tenancy attribute to?

- A. dedicated
- B. isolated
- C. one
- D. reserved

**Answer:** A

**Explanation:**

<http://aws.amazon.com/ec2/dedicated-hosts/>

**NEW QUESTION 97**

Is there any way to own a direct connection to Amazon Web Services?

- A. You can create an encrypted tunnel to VPC, but you don't own the connection.
- B. Yes, it's called Amazon Dedicated Connection.
- C. No, AWS only allows access from the public Internet.
- D. Yes, it's called Direct Connect

**Answer:** D

**NEW QUESTION 101**

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

- A. Yes always
- B. Only if within GovCloud
- C. No
- D. Only if within VPC

**Answer:** A

**NEW QUESTION 102**

To help you manage your Amazon EC2 instances, images, and other Amazon EC2 resources, you can assign your own metadata to each resource in the form of \_\_\_\_\_

- A. special filters
- B. functions
- C. tags
- D. wildcards

**Answer:** C

**NEW QUESTION 106**

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

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 110**

Amazon RDS creates an SSL certificate and installs the certificate on the DB Instance when Amazon RDS provisions the instance. These certificates are signed by a certificate authority. The \_\_\_\_\_ is stored at <https://rds.amazonaws.com/doc/rds-ssl-ca-cert.pem>.

- A. private key
- B. foreign key



- C. public key
- D. protected key

**Answer:** C

**Explanation:**

Amazon RDS creates an SSL certificate and installs the certificate on the DB instance when Amazon RDS provisions the instance. These certificates are signed by a certificate authority. The SSL certificate includes the DB instance endpoint as the Common Name (CN) for the SSL certificate to guard against spoofing attacks. The public key is stored at <https://s3.amazonaws.com/rdsdownloads/rds-combined-ca-bundle.pem>.

**NEW QUESTION 114**

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

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

**Answer:** A

**Explanation:**

The SQL Server bulk copy feature is an efficient means of copying data from a source database to your DB Instance. Bulk copy writes the data that you specify to a data file, such as an ASCII file. You can then run bulk copy again to write the contents of the file to the destination DB Instance.  
<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/SQLServer.Procedural.Importing.html>

**NEW QUESTION 118**

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 119**

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

- A. 1
- B. 50
- C. 20
- D. 10

**Answer:** D

**Explanation:**

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

**NEW QUESTION 123**

If you're unable to connect via SSH to your EC2 instance, which of the following should you check and possibly correct to restore connectivity?

- A. Adjust Security Group to permit egress traffic over TCP port 443 from your IP.
- B. Configure the IAM role to permit changes to security group settings.
- C. Modify the instance security group to allow ingress of ICMP packets from your IP.
- D. Adjust the instance's Security Group to permit ingress traffic over port 22 from your IP.
- E. Apply the most recently released Operating System security patches

**Answer:** D

**Explanation:**

In a VPC everything is allowed out by default. References:

**NEW QUESTION 127**

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

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

**Answer:** C

**Explanation:**

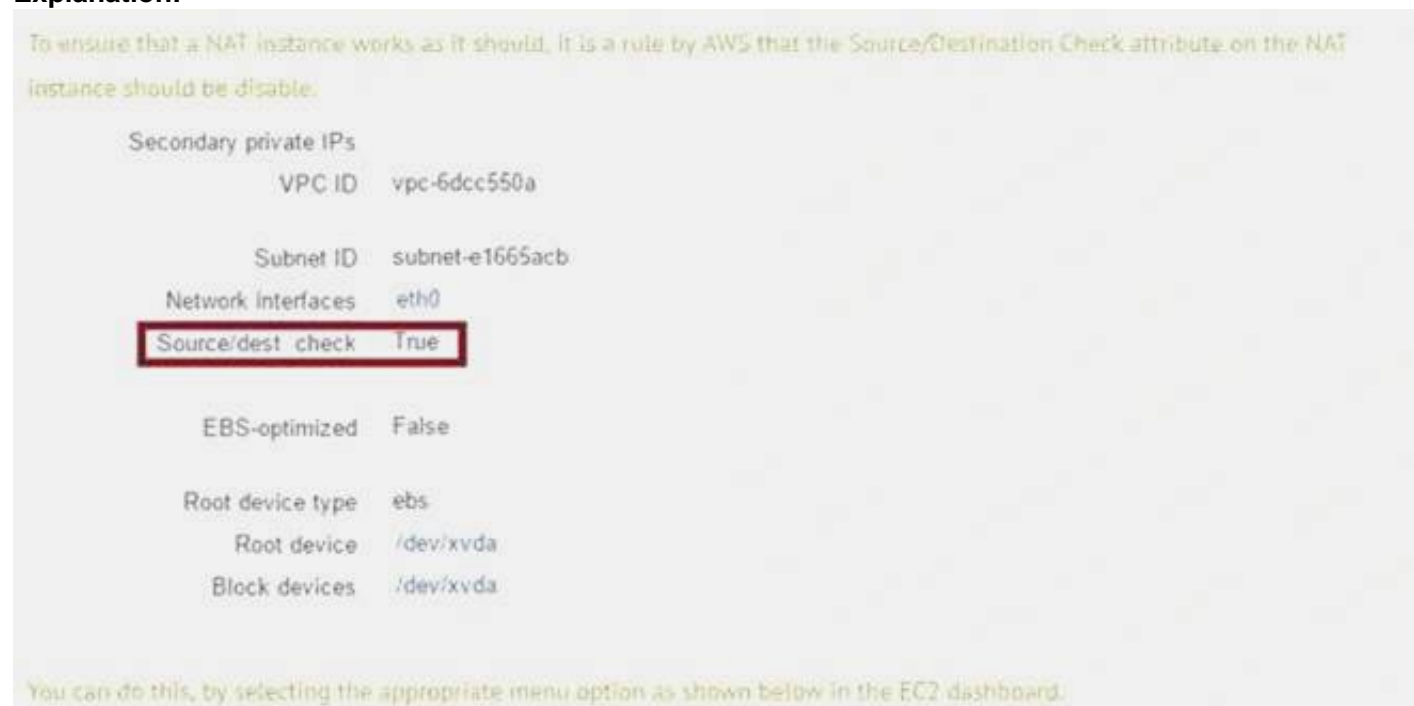
Cross-zone load balancing is always enabled for an Application Load Balancer and is disabled by default for a Classic Load Balancer. If cross-zone load balancing is enabled, the load balancer distributes traffic evenly across all registered instances in all enabled Availability Zones. If cross-zone load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

**NEW QUESTION 131**

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

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

**Answer:** A

**Explanation:****NEW QUESTION 133**

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

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

**Answer:** BDF

**Explanation:**

Enhanced Networking enables you to get significantly higher packet per second (PPS) performance, lower network jitter and lower latencies. This feature uses a new network virtualization stack that provides higher I/O performance and lower CPU utilization compared to traditional implementations. In order to take advantage of Enhanced Networking, you should launch an HVM AMI in VPC, and install the appropriate driver. For instructions on how to enable Enhanced Networking on EC2 instances, see the Enhanced Networking on Linux and Enhanced Networking on Windows tutorials. For availability of this feature by instance, or to learn more, visit the Enhanced Networking FAQ section.

**NEW QUESTION 134**

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

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 139**

Which services allow the customer to retain full administrative privileges of the underlying EC2 instances? (Choose two.)

- A. Amazon Relational Database Service
- B. Amazon Elastic Map Reduce
- C. Amazon ElastiCache
- D. Amazon DynamoDB
- E. AWS Elastic Beanstalk

**Answer:** BE

**NEW QUESTION 140**

Which of the following items are required to allow an application deployed on an EC2 instance to write data to a DynamoDB table? Assume that no security keys are allowed to be stored on the EC2 instance. (Choose two.)

- A. Create an IAM Role that allows write access to the DynamoDB table.
- B. Add an IAM Role to a running EC2 instance.
- C. Create an IAM User that allows write access to the DynamoDB table.
- D. Add an IAM User to a running EC2 instance.
- E. Launch an EC2 Instance with the IAM Role included in the launch configuratio

**Answer:** AB

**NEW QUESTION 142**

You have launched an Amazon Elastic Compute Cloud (EC2) instance into a public subnet with a primary private IP address assigned, an Internet gateway is attached to the VPC, and the public route table is configured to send all Internet-based traffic to the Internet gateway. The instance security group is set to allow all outbound traffic but cannot access the internet. Why is the Internet unreachable from this instance?

- A. The instance does not have a public IP address.
- B. The internet gateway security group must allow all outbound traffic.
- C. The instance security group must allow all inbound traffic.
- D. The instance "Source/Destination check" property must be enable

**Answer:** A

**Explanation:**

Ensure that instances in your subnet have public IP addresses or Elastic IP addresses.  
[https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC\\_Internet\\_Gateway.html](https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Internet_Gateway.html)

**NEW QUESTION 143**

How can the domain's zone apex, for example, "myzoneapexdomain.com", be pointed towards an Elastic Load Balancer?

- A. By using an Amazon Route 53 Alias record
- B. By using an AAAA record
- C. By using an Amazon Route 53 CNAME record
- D. By using an A record

**Answer:** A

**Explanation:**

You can create an alias resource record set at the zone apex. You cannot create a CNAME record at the top node of a DNS namespace, also known as the zone apex. For example, if you register the DNS name example.com, the zone apex is example.com.  
<http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resource-record-sets-choosing-aliasnon-alias.html>

**NEW QUESTION 145**

A client application requires operating system privileges on a relational database server. What is an appropriate configuration for a highly available database architecture?

- A. A standalone Amazon EC2 instance
- B. Amazon RDS in a Multi-AZ configuration
- C. Amazon EC2 instances in a replication configuration utilizing a single Availability Zone
- D. Amazon EC2 instances in a replication configuration utilizing two different Availability Zones

**Answer:** D

**Explanation:**

"A client application requires operating system privileges". You can't have it using RDS. <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html>

**NEW QUESTION 147**

What is a placement group?

- A. A collection of Auto Scaling groups in the same region
- B. A feature that enables EC2 instances to interact with each other via high bandwidth, low latency connections
- C. A collection of authorized CloudFront edge locations for a distribution
- D. A collection of Elastic Load Balancers in the same Region or Availability Zone



**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 Gigabits per second (Gbps) network. Placement groups are recommended for applications that benefit from low network latency, high network throughput, or both.

**NEW QUESTION 152**

You need to configure an Amazon S3 bucket to serve static assets for your public-facing web application. Which methods ensure that all objects uploaded to the bucket are set to public read? (Choose two.)

- A. Set permissions on the object to public read during upload.
- B. Configure the bucket ACL to set all objects to public read.
- C. Configure the bucket policy to set all objects to public read.
- D. Use AWS Identity and Access Management roles to set the bucket to public read.
- E. Amazon S3 objects default to public read, so no action is needed

**Answer:** AC

**Explanation:**

<https://aws.amazon.com/articles/5050>

You can use ACLs to grant permissions to individual AWS accounts; however, it is strongly recommended that you do not grant public access to your bucket using an ACL. So the recommended approach is creating bucket policy, but not ACL. Following link give you an example about how to make the bucket content public.

<http://docs.aws.amazon.com/AmazonS3/latest/dev/HostingWebsiteOnS3Setup.html#step2-addbucket-policy-make-content-public>

**NEW QUESTION 154**

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

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

**Answer:** CE

**Explanation:**

The screenshot shows the AWS documentation for the PUT Object operation. At the top, it asks 'Q: What data consistency model does Amazon S3 employ?'. The answer states: 'Amazon S3 buckets in all Regions provide read-after-write consistency for PUTS of new objects and eventual consistency for overwrite PUTS and DELETES.' Below this, the title 'PUT Object' is shown. Under the 'Description' section, it says: 'This implementation of the PUT operation adds an object to a bucket. You must have WRITE permissions on a bucket to add an object to it.' At the bottom, it states: 'Amazon S3 never adds partial objects: If you receive a success response, Amazon S3 added the entire object to the bucket.'

**NEW QUESTION 159**

A customer needs corporate IT governance and cost oversight of all AWS resources consumed by its divisions. The divisions want to maintain administrative control of the discrete AWS resources they consume and keep those resources separate from the resources of other divisions. Which of the following options, when used together will support the autonomy/control of divisions while enabling corporate IT to maintain governance and cost oversight? (Choose two.)

- A. Use AWS Consolidated Billing and disable AWS root account access for the child accounts.
- B. Enable IAM cross-account access for all corporate IT administrators in each child account.
- C. Create separate VPCs for each division within the corporate IT AWS account.
- D. Use AWS Consolidated Billing to link the divisions' accounts to a parent corporate account.
- E. Write all child AWS CloudTrail and Amazon CloudWatch logs to each child account's Amazon S3 'Log' bucket.

**Answer:** BD

**Explanation:**

B & D are correct when used in combination with each other.

C is theoretically correct by itself, but does not work well with the other choices since it involves only a single AWS account, and the other possibly correct choices (B & D) both involve separate AWS accounts. The question specifically states "Which of the following options, when used together". So C is out.

A is incorrect because you don't want to disable root access to the child accounts (well, except for their access keys for API calls, deleting those is OK).

E is incorrect because it's the exact opposite of a best practice to centralize logs/security audit info across multiple corporate AWS accounts.

[http://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial\\_cross-account-with-roles.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html)

<http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/consolidated-billing.html>

**NEW QUESTION 163**

Which of the following are true regarding AWS CloudTrail? (Choose three.)

- A. CloudTrail is enabled globally
- B. CloudTrail is enabled by default
- C. CloudTrail is enabled on a per-region basis
- D. CloudTrail is enabled on a per-service basis.
- E. Logs can be delivered to a single Amazon S3 bucket for aggregation.
- F. CloudTrail is enabled for all available services within a region.
- G. Logs can only be processed and delivered to the region in which they are generate

**Answer:** ACE

**Explanation:**

A: have a trail with the Apply trail to all regions option enabled.

C: have multiple single region trails.

E: Log files from all the regions can be delivered to a single S3 bucket. Global service events are always delivered to trails that have the Apply trail to all regions option enabled. Events are delivered from a single region to the bucket for the trail. This setting cannot be changed. If you have a single region trail, you should enable the Include global services option. If you have multiple single region trails, you should enable the Include global services option in only one of the trails.

D: Incorrect. Once enabled it is applicable for all the supported services, service can't be selected.

**NEW QUESTION 164**

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

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

**Answer:** A

**Explanation:**

This example looks at using AWS Lambda and Amazon API Gateway to build a dynamic voting application, which receives votes via SMS, aggregates the totals into Amazon DynamoDB, and uses Amazon Simple Storage Service (Amazon S3) to display the results in real time.

<http://www.allthingsdistributed.com/2016/06/aws-lambda-serverless-reference-architectures.html>

**NEW QUESTION 169**

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

- A. Amazon Kinesis
- B. AWS Data Pipeline
- C. Amazon AppStream
- D. Amazon Simple Queue Service

**Answer:** A

**Explanation:**

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

**NEW QUESTION 171**

You have an application running on an Amazon Elastic Compute Cloud instance, that uploads 5 GB video objects to Amazon Simple Storage Service (S3). Video uploads are taking longer than expected, resulting in poor application performance. Which method will help improve performance of your application?

- A. Enable enhanced networking
- B. Use Amazon S3 multipart upload
- C. Leveraging Amazon CloudFront, use the HTTP POST method to reduce latency.
- D. Use Amazon Elastic Block Store Provisioned IOPs and use an Amazon EBS-optimized instance

**Answer:** B

**Explanation:**

Using multipart upload provides the following advantages:

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

**NEW QUESTION 172**

You are designing a web application that stores static assets in an Amazon Simple Storage Service (S3) bucket. You expect this bucket to immediately receive over 150 PUT requests per second. What should you do to ensure optimal performance?

- A. Use multi-part upload.
- B. Add a random prefix to the key names.
- C. Amazon S3 will automatically manage performance at this scale.
- D. Use a predictable naming scheme, such as sequential numbers or date time sequences, in the key names

**Answer:** B

**Explanation:**

If you anticipate that your workload will consistently exceed 100 requests per second, you should avoid sequential key names. If you must use sequential numbers or date and time patterns in key names, add a random prefix to the key name. The randomness of the prefix more evenly distributes key names across multiple index partitions. Examples of introducing randomness are provided later in this topic.

**NEW QUESTION 174**

A company has an AWS account that contains three VPCs (Dev, Test, and Prod) in the same region. Test is peered to both Prod and Dev. All VPCs have non-overlapping CIDR blocks. The company wants to push minor code releases from Dev to Prod to speed up time to market. Which of the following options helps the company accomplish this?

- A. Create a new peering connection Between Prod and Dev along with appropriate routes.
- B. Create a new entry to Prod in the Dev route table using the peering connection as the target.
- C. Attach a second gateway to De
- D. Add a new entry in the Prod route table identifying the gateway as the target.
- E. The VPCs have non-overlapping CIDR blocks in the same accoun
- F. The route tables contain local routes for all VPCs.

**Answer:** A

**NEW QUESTION 176**

You need to pass a custom script to new Amazon Linux instances created in your Auto Scaling group. Which feature allows you to accomplish this?

- A. User data
- B. EC2Config service
- C. IAM roles
- D. AWS Config

**Answer:** A

**Explanation:**

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/user-data.html#user-data-shell-scripts> Not B, because EC2Config is used for Windows instances:  
[http://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/UsingConfig\\_WinAMI.html](http://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/UsingConfig_WinAMI.html)

**NEW QUESTION 180**

Which of the following services natively encrypts data at rest within an AWS region? (Choose two.)

- A. AWS Storage Gateway
- B. Amazon DynamoDB
- C. Amazon CloudFront
- D. Amazon Glacier
- E. Amazon Simple Queue Service

**Answer:** AD

**NEW QUESTION 183**

Which of the following are true regarding encrypted Amazon Elastic Block Store (EBS) volumes? (Choose two.)

- A. Supported on all Amazon EBS volume types
- B. Snapshots are automatically encrypted
- C. Available to all instance types
- D. Existing volumes can be encrypted
- E. shared volumes can be encrypted

**Answer:** AB

**Explanation:**

This feature is supported on all Amazon EBS volume types (General Purpose (SSD), Provisioned IOPS (SSD), and Magnetic). You can access encrypted Amazon EBS volumes the same way you access existing volumes; encryption and decryption are handled transparently and they require no additional action from you, your Amazon EC2 instance, or your application. Snapshots of encrypted Amazon EBS volumes are automatically encrypted, and volumes that are created from encrypted Amazon EBS snapshots are also automatically encrypted.

**NEW QUESTION 186**

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

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

**Answer:** B

**Explanation:**

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



**NEW QUESTION 190**

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

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

**Answer: D**

**Explanation:**

Disabling Source/Destination Checks.

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

**NEW QUESTION 192**

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

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

**Answer: C**

**Explanation:**

Geolocation routing lets you choose the resources that serve your traffic based on the geographic location of your users, meaning the location from which DNS queries originate. For example, you might want all queries from Africa to be routed to a web server with an IP address of 192.0.2.111. Another possible use is for balancing load across endpoints in a predictable, easy-to-manage way, so that each user location is consistently routed to the same endpoint.

<http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html#routing-policyweighted>

**NEW QUESTION 194**

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

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

**Answer: A**

**NEW QUESTION 196**

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

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

**Answer: B**

**NEW QUESTION 197**

By definition a public subnet within a VPC is one that

- A. In its routing table it has at least one route that uses an Internet Gateway (IGW).
- B. Has at least one route in its routing table that routes via a Network Address Translation (NAT) instance.
- C. Where the Network Access Control List (NACL) permitting outbound traffic to 0.0.0.0/0.
- D. Has had the public subnet check box ticked when setting up this subnet in the VPC consol

**Answer: A**

**NEW QUESTION 200**

You work for a toy company that has a busy online store. As you are approaching christmas you find that your store is getting more and more traffic. You ensure

that the web tier of your store is behind an Auto Scaling group, however you notice that the web tier is frequently scaling, sometimes multiple times in an hour, only to scale back after peak usage. You need to prevent this so that Auto Scaling does not scale as rapidly, just to scale back again. What option would help you to achieve this?

- A. Configure Auto Scaling to terminate your oldest instances first, then adjust your CloudWatch alarm.
- B. Configure Auto Scaling to terminate your newest instances first, then adjust your CloudWatch alarm.
- C. Change your Auto Scaling so that it only scales at scheduled times.
- D. Modify the Auto Scaling group cool-down timers & modify the Amazon CloudWatch alarm period that triggers your Auto Scaling scale down policy.

**Answer:** D

**Explanation:**

### Auto Scaling Cooldowns

The Auto Scaling cooldown period is a configurable setting for your Auto Scaling group that helps to ensure that **Auto Scaling doesn't launch or terminate additional instances before the previous scaling activity takes effect.** After the Auto Scaling group dynamically scales using a simple scaling policy, Auto Scaling waits for the cooldown period to complete before resuming scaling activities. When you manually scale your Auto Scaling group, the default is not to wait for the cooldown period, but you can override the default and honor the cooldown period. Note that if an instance becomes unhealthy, Auto Scaling does not wait for the cooldown period to complete before replacing the unhealthy instance.

#### NEW QUESTION 204

You are a student currently learning about the different AWS services. Your employer asks you to tell him a bit about Amazon's glacier service. Which of the following best describes the use cases for Glacier?

- A. Infrequently accessed data & data archives
- B. Hosting active databases
- C. Replicating Files across multiple availability zones and regions
- D. Frequently Accessed Data

**Answer:** A

#### NEW QUESTION 205

You have uploaded a file to S3. What HTTP code would indicate that the upload was successful?

- A. HTTP 404
- B. HTTP 501
- C. HTTP 200
- D. HTTP 307

**Answer:** C

#### NEW QUESTION 207

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

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

**Answer:** E

#### NEW QUESTION 209

You can select a specific Availability Zone in which to place your DynamoDB Table

- A. True
- B. False

**Answer:** B

#### NEW QUESTION 212

To retrieve instance metadata or userdata you will need to use the following IP Address;

- A. <http://127.0.0.1>
- B. <http://192.168.0.254>
- C. <http://10.0.0.1>
- D. <http://169.254.169.254>

**Answer:** D

**NEW QUESTION 215**

When you create new subnets within a custom VPC, by default they can communicate with each other, across availability zones.

- A. True
- B. False

**Answer:** A

**NEW QUESTION 216**

Which of the services below do you get root access to?

- A. Elasticache & Elastic MapReduce
- B. RDS & DynamoDB
- C. EC2 & Elastic MapReduce
- D. Elasticache & DynamoDB

**Answer:** C

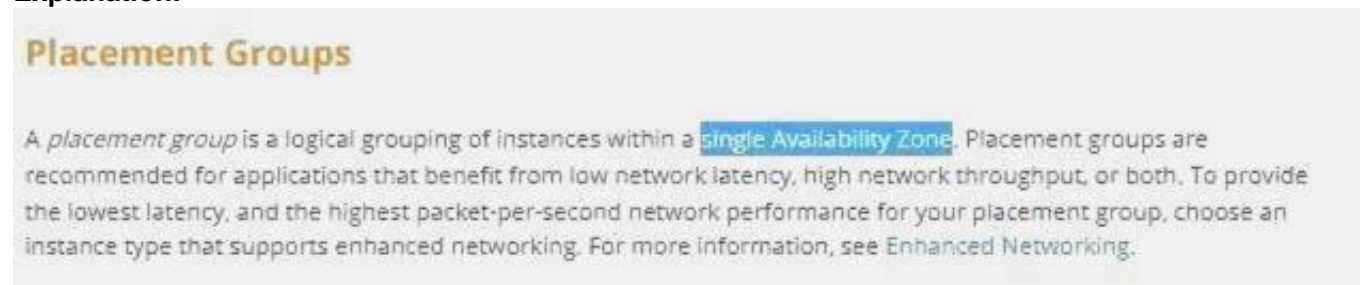
**NEW QUESTION 217**

Placement Groups can be created across 2 or more Availability Zones.

- A. True
- B. False

**Answer:** B

**Explanation:**



**NEW QUESTION 219**

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

- A. True
- B. False

**Answer:** B

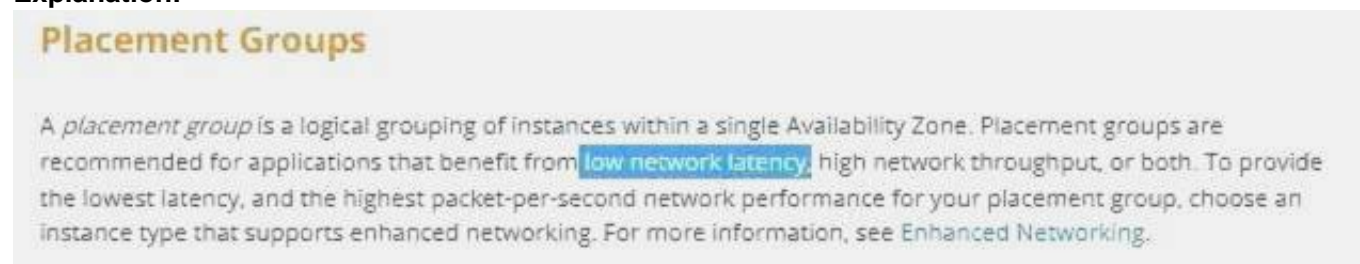
**NEW QUESTION 222**

You have a high performance compute application and you need to minimize network latency between EC2 instances as much as possible. What can you do to achieve this?

- A. Use Elastic Load Balancing to load balance traffic between availability zones
- B. Create a CloudFront distribution and to cache objects from an S3 bucket at Edge Locations.
- C. Create a placement group within an Availability Zone and place the EC2 instances within that placement group.
- D. Deploy your EC2 instances within the same region, but in different subnets and different availability zones so as to maximize redundancy.

**Answer:** C

**Explanation:**



**NEW QUESTION 225**

You are hosting a website in Ireland called aloud.guru and you decide to have a static DR site available on S3 in the event that your primary site would go down. Your bucket name is also called "acloudguru". What would be the S3 URL of the static website?

- A. <https://acloudguru.s3-website-eu-west-1.amazonaws.com>
- B. <https://s3-eu-east-1.amazonaws.com/acloudguru>
- C. <https://acloudguru.s3-website-us-east-1.amazonaws.com>
- D. <https://s3-eu-central-1.amazonaws.com/acloudguru>



**Answer:** A

#### NEW QUESTION 229

When trying to grant an amazon account access to S3 using access control lists what method of identification should you use to identify that account with?

- A. The email address of the account or the canonical user ID
- B. The AWS account number
- C. The ARN
- D. An email address with a 2FA token

**Answer:** A

#### NEW QUESTION 234

You are a solutions architect working for a company that specializes in ingesting large data feeds (using Kinesis) and then analyzing these feeds using Elastic Map Reduce (EMR). The results are then stored on a custom MySQL database which is hosted on an EC2 instance which has 3 volumes, the root/boot volume, and then 2 additional volumes which are striped in to a RAID 1. Your company recently had an outage and lost some key data and have since decided that they will need to run nightly back ups. Your application is only used during office hours, so you can afford to have some down time in the middle of the night if required. You decide to take a snapshot of all three volumes every 24 hours. In what manner should you do this?

- A. Take a snapshot of each volume independently, while the EC2 instance is running.
- B. Stop the EC2 instance and take a snapshot of each EC2 instance independentl
- C. Once the snapshots are complete, start the EC2 instance and ensure that all relevant volumes are remounted.
- D. Add two additional volumes to the existing RAID 0 volume and mirror these volumes creating a RAID 10. Take a snap of only the two new volumes.
- E. Create a read replica of the existing EC2 instance and then take your snapshots from the read replica and not the live EC2 instance.

**Answer:** B

#### NEW QUESTION 237

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

- A. True
- B. False

**Answer:** A

#### NEW QUESTION 240

What is a Security Group?

- A. None of these.
- B. A list of users that can access Amazon EC2 instances.
- C. An Access Control List (ACL) for AWS resources.
- D. It acts as a virtual firewall that controls the traffic for one or more instance

**Answer:** D

#### NEW QUESTION 241

While performing volume status checks using volume status checks, if the status is insufficient-data, what does it mean?

- A. checks may still be in progress on the volume
- B. check has passed
- C. check has failed
- D. there is no such status

**Answer:** A

#### Explanation:

Volume status checks are automated tests that run every 5 minutes and return a pass or fail status. If all checks pass, the status of the volume is ok. If a check fails, the status of the volume is impaired. If the status is insufficient-data, the checks may still be in progress on the volume.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/monitoring-volume-status.html>

#### NEW QUESTION 245

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

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

**Answer:** A

#### Explanation:

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

**NEW QUESTION 249**

You have a business-critical two-tier web app currently deployed in two AZs in a single region, using Elastic Load Balancing and Auto Scaling. The app depends on synchronous replication (very low latency connectivity) at the database layer. The application needs to remain fully available even if one application AZ goes off-line, and Auto Scaling cannot launch new instances in the remaining Availability Zones. How can the current architecture be enhanced to ensure this?

- A. Deploy in two regions using Weighted Round Robin (WRR), with Auto Scaling minimums set for 50 percent peak load per Region.
- B. Deploy in two regions using Weighted Round Robin (WRR), with Auto Scaling minimums set for 100 percent peak load per region.
- C. Deploy in three Availability Zones, with Auto Scaling minimum set to handle 50 percent peak load per zone.
- D. Deploy in three Availability Zones, with Auto Scaling minimum set to handle 33 percent peak load per zone.

**Answer:** C

**NEW QUESTION 253**

You are developing a highly available web application using stateless web servers. Which services are suitable for storing session state data? (Choose three.)

- A. Amazon CloudWatch
- B. Amazon Relational Database Service (RDS)
- C. Elastic Load Balancing
- D. Amazon ElastiCache
- E. AWS Storage Gateway
- F. Amazon DynamoDB

**Answer:** BDF

**NEW QUESTION 256**

Which of the following requires a custom CloudWatch metric to monitor?

- A. Memory use
- B. CPU use
- C. Disk read operations
- D. Network in
- E. Estimated charges

**Answer:** A

**NEW QUESTION 260**

What action is required to establish a VPC VPN connection between an on-premises data center and an Amazon VPC virtual private gateway?

- A. Modify the main route table to allow traffic to a network address translation instance.
- B. Use a dedicated network address translation instance in the public subnet.
- C. Assign a static Internet-routable IP address to an Amazon VPC customer gateway.
- D. Establish a dedicated networking connection using AWS Direct Connect.

**Answer:** C

**NEW QUESTION 262**

What is the minimum size of an EBS volume as per AWS?

- A. 2TB
- B. 1GiB
- C. 1GB
- D. 1Byte

**Answer:** B

**NEW QUESTION 263**

Amazon RDS provides a facility to modify the back-up retention policy for automated backups, with a value of 0 indicating for no backup retention. What is the maximum retention period allowed in days?

- A. 45
- B. 35
- C. 15
- D. 10

**Answer:** B

**Explanation:**

**NEW QUESTION 264**

Your existing web application requires a persistent key-value store database that must service 50,000 reads/second. Your company is looking at 10% growth in traffic and data volume month over month for the next several years. Which service meets these requirements?

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

**Answer: D**

**NEW QUESTION 269**

An application on an Amazon EC2 instance routinely stops responding to requests and requires a reboot to recover. The application logs are already exported into Amazon CloudWatch, and you notice that the problem consistently follows the appearance of a specific message in the log. The application team is working to address the bug, but has not provided a date for the fix. What workaround can you implement to automate recovery of the instance until the fix is deployed?

- A. Create an Amazon CloudWatch alarm on an Amazon CloudWatch Logs filter for that message; based on that alarm, trigger an Amazon CloudWatch action to reboot the instance.
- B. Create an AWS CloudTrail alarm on low CPU; based on that alarm, trigger an Amazon SNS message to the Operations team.
- C. Create an Amazon CloudWatch alarm on instance memory usage; based on that alarm, trigger an Amazon CloudWatch action to reboot the instance.
- D. Create an AWS CloudTrail alarm to detect the deadlock; based on that alarm, trigger an Amazon SNS message to the Operations team.

**Answer: C**

**NEW QUESTION 271**

Your Amazon VPC has a public subnet with a route that sends all Internet traffic to the Internet gateway. An Amazon EC2 instance in the public subnet has an assigned private IP address. The instance belongs to a security group set to allow all outbound traffic. The instance cannot access the Internet. Why could the Internet be unreachable from this instance?

- A. The instance does not have a public IP address.
- B. The internet gateway security group must allow all outbound traffic.
- C. The instance security group must allow all inbound traffic.
- D. The instance "Source/Destination check" property must be enable

**Answer: A**

**NEW QUESTION 274**

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

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

**Answer: AD**

**NEW QUESTION 277**

What services will help identify Amazon EC2 instances with underutilized CPU capacity? (Select TWO.)

- A. Amazon CloudWatch
- B. Cost Explorer
- C. AWS Trusted Advisor
- D. AWS CloudTrail
- E. Amazon EC2 usage reports

**Answer:** AE

**NEW QUESTION 279**

You have set up an Auto Scaling group. The cool down period for the Auto Scaling group is 7 minutes. The first instance is launched after 3 minutes, while the second instance is launched after 4 minutes.

How many minutes after the first instance is launched will Auto Scaling accept another scaling activity request?

- A. 11 minutes
- B. 7 minutes
- C. 10 minutes
- D. 14 minutes

**Answer:** A

**Explanation:**

If an Auto Scaling group is launching more than one instance, the cool down period for each instance starts after that instance is launched. The group remains locked until the last instance that was launched has completed its cool down period. In this case the cool down period for the first instance starts after 3 minutes and finishes at the 10th minute (3+7 cool down), while for the second instance it starts at the 4th minute and finishes at the 11th minute (4+7 cool down). Thus, the Auto Scaling group will receive another request only after 11 minutes.

**NEW QUESTION 284**

In the context of AWS support, why must an EC2 instance be unreachable for 20 minutes rather than allowing customers to open tickets immediately?

- A. Because most reachability issues are resolved by automated processes in less than 20 minutes
- B. Because all EC2 instances are unreachable for 20 minutes every day when AWS does routine maintenance
- C. Because all EC2 instances are unreachable for 20 minutes when first launched
- D. Because of all the reasons listed here

**Answer:** A

**Explanation:**

An EC2 instance must be unreachable for 20 minutes before opening a ticket, because most reachability issues are resolved by automated processes in less than 20 minutes and will not require any action on the part of the customer. If the instance is still unreachable after this time frame has passed, then you should open a case with support.

**NEW QUESTION 288**

Do Amazon EBS volumes persist independently from the running life of an Amazon EC2 instance?

- A. Yes, they do but only if they are detached from the instance.
- B. No, you cannot attach EBS volumes to an instance.
- C. No, they are dependent.
- D. Yes, they do.

**Answer:** D

**Explanation:**

An Amazon EBS volume behaves like a raw, unformatted, external block device that you can attach to a single instance. The volume persists independently from the running life of an Amazon EC2 instance.

**NEW QUESTION 292**

A user has created an EBS volume with 1000 IOPS. What is the average IOPS that the user will get for most of the year as per EC2 SLA if the instance is attached to the EBS optimized instance?

- A. 950
- B. 990
- C. 1000
- D. 900

**Answer:** D

**Explanation:**

As per AWS SLA if the instance is attached to an EBS-Optimized instance, then the Provisioned IOPS volumes are designed to deliver within 10% of the provisioned IOPS performance 99.9% of the time in a given year. Thus, if the user has created a volume of 1000 IOPS, the user will get a minimum 900 IOPS 99.9% time of the year.

**NEW QUESTION 296**

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. References:

**NEW QUESTION 300**

You are in the process of creating a Route 53 DNS failover to direct traffic to two EC2 zones. Obviously, if one fails, you would like to direct Route 53 traffic to the other region. Each region has an ELB with some instances being distributed. What is the best way for you to configure the Route 53 health check?

- A. Route 53 doesn't support ELB with an internal health check
- B. You need to create your own Route 53 health check of the ELB
- C. Route 53 natively supports ELB with an internal health check
- D. Turn "Evaluate target health" off and "Associate with Health Check" on and R53 will use the ELB's internal health check.
- E. Route 53 doesn't support ELB with an internal health check
- F. You need to associate your resource record set for the ELB with your own health check
- G. Route 53 natively supports ELB with an internal health check
- H. Turn "Evaluate target health" on and "Associate with Health Check" off and R53 will use the ELB's internal health check.

**Answer:** D

**Explanation:**

With DNS Failover, Amazon Route 53 can help detect an outage of your website and redirect your end users to alternate locations where your application is operating properly. When you enable this feature, Route 53 uses health checks--regularly making Internet requests to your application's endpoints from multiple locations around the world--to determine whether each endpoint of your application is up or down. To enable DNS Failover for an ELB endpoint, create an Alias record pointing to the ELB and set the "Evaluate Target Health" parameter to true. Route 53 creates and manages the health checks for your ELB automatically. You do not need to create your own Route 53 health check of the ELB. You also do not need to associate your resource record set for the ELB with your own health check, because Route 53 automatically associates it with the health checks that Route 53 manages on your behalf. The ELB health check will also inherit the health of your backend instances behind that ELB.

**NEW QUESTION 302**

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

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

**Answer:** C

**Explanation:**

Auto Scaling can start and stop the instance at a pre-defined time. Here, the total running time is unknown. Thus, the user has to use the CloudWatch alarm, which monitors the CPU utilization.

The user can create an alarm that is triggered when the average CPU utilization percentage has been lower than 10 percent for 24 hours, signaling that it is idle and no longer in use. When the utilization is below the threshold limit, it will terminate the instance as a part of the instance action. References:

**NEW QUESTION 305**

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

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

**Answer:** B

**Explanation:**

By default, all AWS accounts are limited to 5 Elastic IP addresses per region for each AWS account, because public (IPv4) Internet addresses are a scarce public resource. AWS strongly encourages you to use an EIP primarily for load balancing use cases, and use DNS hostnames for all other inter-node communication. If you feel your architecture warrants additional EIPs, you would need to complete the Amazon EC2 Elastic IP Address Request Form and give reasons as to your need for additional addresses. References:

**NEW QUESTION 307**

While using the EC2 GET requests as URLs, the is the URL that serves as the entry point for the web service.

- A. token
- B. endpoint
- C. action
- D. None of these

**Answer:** B

**Explanation:**

The endpoint is the URL that serves as the entry point for the web service.

**NEW QUESTION 309**

You have been given a scope to deploy some AWS infrastructure for a large organization. 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.

**NEW QUESTION 314**

A major finance organization has engaged your company to set up a large data mining application. Using AWS you decide the best service for this is Amazon Elastic MapReduce(EMR) which you know uses Hadoop. Which of the following statements best describes Hadoop?

- A. Hadoop is 3rd Party software which can be installed using AMI
- B. Hadoop is an open source python web framework
- C. Hadoop is an open source Java software framework
- D. Hadoop is an open source JavaScript framework

**Answer: C**

**Explanation:**

Amazon EMR uses Apache Hadoop as its distributed data processing engine. Hadoop is an open source, Java software framework that supports data-intensive distributed applications running on large clusters of commodity hardware. Hadoop implements a programming model named "MapReduce," where the data is divided into many small fragments of work, each of which may be executed on any node in the cluster.

This framework has been widely used by developers, enterprises and startups and has proven to be a reliable software platform for processing up to petabytes of data on clusters of thousands of commodity machines.

**NEW QUESTION 319**

\_\_\_\_\_ is a fast, flexible, 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, flexible, 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.

**NEW QUESTION 323**

You need to set up a complex network infrastructure for your organization that will be reasonably easy to deploy, replicate, control, and track changes on. Which AWS service would be best to use to help you accomplish this?

- A. AWS Import/Export
- B. AWS CloudFormation
- C. Amazon Route 53
- D. Amazon CloudWatch

**Answer: B**

**Explanation:**

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. You don't need to individually create and configure AWS resources and figure out what's dependent on what. AWS CloudFormation handles all of that.

**NEW QUESTION 324**

In an experiment, if the minimum size for an Auto Scaling group is 1 instance, which of the following statements holds true when you terminate the running instance?

- A. Auto Scaling must launch a new instance to replace it.
- B. Auto Scaling will raise an alarm and send a notification to the user for action.
- C. Auto Scaling must configure the schedule activity that terminates the instance after 5 days.
- D. Auto Scaling will terminate the experiment

**Answer: A**

**Explanation:**

If the minimum size for an Auto Scaling group is 1 instance, when you terminate the running instance, Auto Scaling must launch a new instance to replace it.

**NEW QUESTION 327**

Your company has been storing a lot of data in Amazon Glacier and has asked for an inventory of what is in there exactly. So you have decided that you need to download a vault inventory. Which of the following statements is incorrect in relation to Vault Operations in Amazon Glacier?

- A. You can use Amazon Simple Notification Service (Amazon SNS) notifications to notify you when the job completes.
- B. A vault inventory refers to the list of archives in a vault.
- C. You can use Amazon Simple Queue Service (Amazon SQS) notifications to notify you when the job completes.
- D. Downloading a vault inventory is an asynchronous operation.

**Answer: C**

**Explanation:**

Amazon Glacier supports various vault operations.

A vault inventory refers to the list of archives in a vault. For each archive in the list, the inventory provides archive information such as archive ID, creation date, and size. Amazon Glacier updates the vault inventory approximately once a day, starting on the day the first archive is uploaded to the vault. A vault inventory must exist for you to be able to download it.

Downloading a vault inventory is an asynchronous operation. You must first initiate a job to download the inventory. After receiving the job request, Amazon Glacier prepares your inventory for download. After the job completes, you can download the inventory data.

Given the asynchronous nature of the job, you can use Amazon Simple Notification Service (Amazon SNS) notifications to notify you when the job completes. You can specify an Amazon SNS topic for each individual job request or configure your vault to send a notification when specific vault events occur. Amazon Glacier prepares an inventory for each vault periodically, every 24 hours. If there have been no archive additions or deletions to the vault since the last inventory, the inventory date is not updated. When you initiate a job for a vault inventory, Amazon Glacier returns the last inventory it generated, which is a point-in-time snapshot and not real-time data. You might not find it useful to retrieve vault inventory for each archive upload. However, suppose you maintain a database on the client-side associating metadata about the archives you upload to Amazon Glacier. Then, you might find the vault inventory useful to reconcile information in your database with the actual vault inventory.

**NEW QUESTION 332**

What does the following policy for Amazon EC2 do?

```
{
  "Statement": [{
    "Effect": "Allow",
    "Action": "ec2:Describe*",
    "Resource": "*"
  }]
}
```

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

**Answer: A**

**Explanation:**

You can use IAM policies to control the actions that your users can perform against your EC2 resources. For instance, a policy with the following statement will allow users to perform actions whose name start with "Describe" against all your EC2 resources.

```
{
  "Statement": [{
    "Effect": "Allow",
    "Action": "ec2:Describe*",
    "Resource": "*"
  }]
}
```

**NEW QUESTION 334**

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

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

**Answer: A**

**Explanation:**

You can set up a variety of failover configurations using Amazon Route 53 alias: weighted, latency, geolocation routing, and failover resource record sets.

Active-active failover: Use this failover configuration when you want all of your resources to be available the majority of the time. When a resource becomes unavailable, Amazon Route 53 can detect that it's unhealthy and stop including it when responding to queries. Active-passive failover:

Use this failover configuration when you want a primary group of resources to be available the majority of the time and you want a secondary group of resources to be on standby in case all of the primary resources become unavailable. When responding to queries, Amazon Route 53 includes only the healthy primary resources. If all of the primary resources are unhealthy, Amazon Route 53 begins to include only the healthy secondary resources in response to DNS queries.

Active-active-passive and other mixed configurations: You can combine alias and non-alias resource record sets to produce a variety of Amazon Route 53

behaviors.

**NEW QUESTION 338**

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"
      }
    }
  }
}
```

**NEW QUESTION 339**

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

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

**Answer:** B

**Explanation:**

Amazon S3 supports several mechanisms that give you 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 342**

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

- A. The IP address may be attached to one of the instances



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

**Answer:** D

**Explanation:**

A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. When the user is launching an instance he needs to select an option which attaches a public IP to the instance. If the user has not selected the option to attach the public IP, then it will only have a private IP when launched. If the user wants to connect to an instance from the Internet, he should create an elastic IP with VPC. If the elastic IP is a part of EC2 Classic, it cannot be assigned to a VPC instance.

**NEW QUESTION 344**

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.

**NEW QUESTION 347**

Which one of the following answers is not a possible state of Amazon CloudWatch Alarm?

- A. INSUFFICIENT\_DATA
- B. ALARM
- C. OK
- D. STATUS\_CHECK\_FAILED

**Answer:** D

**Explanation:**

Amazon CloudWatch Alarms have three possible states: OK: The metric is within the defined threshold ALARM: The metric is outside of the defined threshold INSUFFICIENT\_DATA: The alarm has just started, the metric is not available, or not enough data is available for the metric to determine the alarm state

**NEW QUESTION 352**

Which of the following strategies can be used to control access to your Amazon EC2 instances?

- A. DB security groups
- B. IAM policies
- C. None of these
- D. EC2 security groups

**Answer:** D

**Explanation:**

IAM policies allow you to specify what actions your IAM users are allowed to perform against your EC2 Instances. However, when it comes to access control, security groups are what you need in order to define and control the way you want your instances to be accessed, and whether or not certain kind of communications are allowed or not.

**NEW QUESTION 354**

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

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

**Answer:** D

**Explanation:**

Amazon EC2 supports the following storage options: Amazon Elastic Block Store (Amazon EBS) Amazon EC2 Instance Store Amazon Simple Storage Service (Amazon S3)

**NEW QUESTION 355**

You are signed in as root user on your account but there is an Amazon S3 bucket under your account that you cannot access. What is a possible reason for this?

- A. An IAM user assigned a bucket policy to an Amazon S3 bucket and didn't specify the root user as a principal
- B. The S3 bucket is full.
- C. The S3 bucket has reached the maximum number of objects allowed.
- D. You are in the wrong availability zone

**Answer:** A

**Explanation:**

With IAM, you can centrally manage users, security credentials such as access keys, and permissions that control which AWS resources users can access. In some cases, you might have an IAM user with full access to IAM and Amazon S3. If the IAM user assigns a bucket policy to an Amazon S3 bucket and doesn't specify the root user as a principal, the root user is denied access to that bucket. However, as the root user, you can still access the bucket by modifying the bucket policy to allow root user access.

**NEW QUESTION 358**

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

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

**Answer:** A

**Explanation:**

DynamoDB supports two types of secondary indexes:

Local secondary index --an index that has the same hash key as the table, but a different range key. A local secondary index is "local" in the sense that every partition of a local secondary index is scoped to a table partition that has the same hash key.

Global secondary index --an index with a hash and range key that can be different from those on the table. A global secondary index is considered "global" because queries on the index can span all of the data in a table, across all partitions.

**NEW QUESTION 361**

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

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

**Answer:** B

**Explanation:**

Amazon VPC provides two features that you can use to increase security for your VPC:

Security groups--Act as a firewall for associated Amazon EC2 instances, controlling both inbound and outbound traffic at the instance level.

Network access control lists (ACLs)--Act as a firewall for associated subnets, controlling both inbound and outbound traffic at the subnet level.

Security groups are stateful: (Return traffic is automatically allowed, regardless of any rules) Network ACLs are stateless: (Return traffic must be explicitly allowed by rules)

**NEW QUESTION 365**

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 368**

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

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 369**

Which one of the below is not an AWS Storage Service?

- A. Amazon S3
- B. Amazon Glacier
- C. Amazon CloudFront
- D. Amazon EBS

**Answer:** C

**Explanation:**

AWS Storage Services are: Amazon S3  
Amazon Glacier Amazon EBS  
AWS Storage Gateway

**NEW QUESTION 371**

You are very concerned about security on your network because you have multiple programmers testing APIs and SDKs and you have no idea what is happening. You think CloudTrail may help but are not sure what it does. Which of the following statements best describes the AWS service CloudTrail?

- A. With AWS CloudTrail you can get a history of AWS API calls and related events for your account.
- B. With AWS CloudTrail you can get a history of IAM users for your account.
- C. With AWS CloudTrail you can get a history of S3 logfiles for your account.
- D. With AWS CloudTrail you can get a history of CloudFormation JSON scripts used for your account.

**Answer:** A

**Explanation:**

With AWS CloudTrail, you can get a history of AWS API calls for your account, including API calls made via the AWS Management Console, the AWS SDKs, the command line tools, and higher-level AWS services. You can also identify which users and accounts called AWS APIs for services that support CloudTrail, the source IP address the calls were made from, and when the calls occurred. You can identify which users and accounts called AWS for services that support CloudTrail, the source IP address the calls were made from, and when the calls occurred. You can integrate CloudTrail into applications using the API, automate trail creation for your organization, check the status of your trails, and control how administrators turn CloudTrail logging on and off.

**NEW QUESTION 375**

A user is currently building a website which will require a large number of instances in six months, when a demonstration of the new site will be given upon launch. Which of the below mentioned options allows the user to procure the resources beforehand so that they need not worry about infrastructure availability during the demonstration?

- A. Procure all the instances as reserved instances beforehand.
- B. Launch all the instances as part of the cluster group to ensure resource availability.
- C. Pre-warm all the instances one month prior to ensure resource availability.
- D. Ask AWS now to procure the dedicated instances in 6 month

**Answer:** A

**Explanation:**

Amazon Web Services has massive hardware resources at its data centers, but they are finite. The best way for users to maximize their access to these resources is by reserving a portion of the computing capacity that they require. This can be done through reserved instances. With reserved instances, the user literally reserves the computing capacity in the Amazon Web Services cloud.

**NEW QUESTION 376**

.....

## Thank You for Trying Our Product

\* 100% Pass or Money Back

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

\* One year free update

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

\* Trusted by Millions

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

\* Shop Securely

All transactions are protected by VeriSign!

**100% Pass Your SAA-C01 Exam with Our Prep Materials Via below:**

<https://www.certleader.com/SAA-C01-dumps.html>