

## SAA-C01 Dumps

### AWS Certified Solutions Architect - Associate

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



**NEW QUESTION 1**

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

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

**Answer: B**

**NEW QUESTION 2**

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 3**

Your company runs a customer facing event registration site. This site is built with a 3-tier architecture with web and application tier servers and a MySQL database. The application requires 6 web tier servers and 6 application tier servers for normal operation, but can run on a minimum of 65% server capacity and a single MySQL database. When deploying this application in a region with three availability zones (AZs) which architecture provides high availability?

- A. A web tier deployed across 2 AZs with 3 EC2 (Elastic Compute Cloud) instances in each AZ inside an Auto Scaling Group behind an ELB (elastic load balancer), and an application tier deployed across 2 AZs with 3 EC2 instances in each AZ inside an Auto Scaling Group behind an ELB, and one RDS (Relational Database Service) instance deployed with read replicas in the other AZ.
- B. A web tier deployed across 3 AZs with 2 EC2 (Elastic Compute Cloud) instances in each AZ inside an Auto Scaling Group behind an ELB (elastic load balancer) and an application tier deployed across 3 AZs with 2 EC2 instances in each AZ inside an Auto Scaling Group behind an ELB and one RDS (Relational Database Service) Instance deployed with read replicas in the two other AZs.
- C. A web tier deployed across 2 AZs with 3 EC2 (Elastic Compute Cloud) instances in each AZ inside an Auto Scaling Group behind an ELB (elastic load balancer) and an application tier deployed across 2 AZs with 3 EC2 instances in each AZ inside an Auto Scaling Group behind an ELB and a Multi-AZ RDS (Relational Database Service) deployment.
- D. A web tier deployed across 3 AZs with 2 EC2 (Elastic Compute Cloud) instances in each AZ inside an Auto Scaling Group behind an ELB (elastic load balancer). And an application tier deployed across 3 AZs with 2 EC2 instances in each AZ inside an Auto Scaling Group behind an ELB.
- E. And a Multi-AZ RDS (Relational Database services) deployment.

**Answer: D**

**Explanation:**

Amazon RDS Multi-AZ Deployments

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

Enhanced Durability

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

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

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

Increased Availability

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

The availability benefits of Multi-AZ deployments also extend to planned maintenance and backups. In the case of system upgrades like OS patching or DB Instance scaling, these operations are applied first on the standby, prior to the automatic failover. As a result, your availability impact is, again, only the time required for automatic failover to complete.

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

On instance failure in Amazon Aurora deployments, Amazon RDS uses RDS Multi-AZ technology to automate failover to one of up to 15 Amazon Aurora Replicas

you have created in any of three Availability Zones. If no Amazon Aurora Replicas have been provisioned, in the case of a failure, Amazon RDS will attempt to create a new Amazon Aurora DB instance for you automatically. <https://www.airpair.com/aws/posts/building-a-scalable-web-app-on-amazon-web-services-p1>

**NEW QUESTION 4**

You are responsible for a legacy web application whose server environment is approaching end of life. You would like to migrate this application to AWS as quickly as possible, since the application environment currently has the following limitations:

The VM's single 10GB VMDK is almost full.

The virtual network interface still uses the 10Mbps driver, which leaves your 100Mbps WAN connection completely underutilized.

It is currently running on a highly customized Windows VM within a VMware environment. You do not have the installation media.

This is a mission-critical application with an RTO (Recovery Time Objective) of 8 hours. RPO (Recovery Point Objective) of 1 hour. How could you best migrate this application to AWS while meeting your business continuity requirements?

- A. Use the EC2 VM Import Connector for vCenter to import the VM into EC2.
- B. Use Import/Export to import the VM as an ESS snapshot and attach to EC2.
- C. Use S3 to create a backup of the VM and restore the data into EC2.
- D. Use the ec2-bundle-instance API to import an image of the VM into EC2.

**Answer:** A

**Explanation:**

<https://aws.amazon.com/developertools/2759763385083070>

**NEW QUESTION 5**

A customer has established an AWS Direct Connect connection to AWS. The link is up and routes are being advertised from the customer's end, however the customer is unable to connect from EC2 instances inside its VPC to servers residing in its datacenter.

Which of the following options provide a viable solution to remedy this situation? (Choose two.)

- A. Add a route to the route table with an iPsec VPN connection as the target.
- B. Enable route propagation to the virtual pinnate gateway (VGW).
- C. Enable route propagation to the customer gateway (CGW).
- D. Modify the route table of all Instances using the 'route' command.
- E. Modify the Instances VPC subnet route table by adding a route back to the customer's on-premises environment.

**Answer:** BE

**NEW QUESTION 6**

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

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

**Answer:** C

**Explanation:**

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

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

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

**NEW QUESTION 7**

You are designing a multi-platform web application for AWS. The application will run on EC2 instances and will be accessed from PCs, tablets and smart phones. Supported accessing platforms are Windows, MacOS, IOS and Android. Separate sticky session and SSL certificate setups are required for different platform types. Which of the following describes the most cost-effective and performance-efficient architecture setup?

- A. Setup a hybrid architecture to handle session state and SSL certificates on-prem and separate EC2 Instance groups running web applications for different platform types running in a VPC.
- B. Set up one ELB for all platforms to distribute load among multiple instances under it. Each EC2 instance implements all functionality for a particular platform.
- C. Set up two ELBs. The first ELB handles SSL certificates for all platforms and the second ELB handles session stickiness for all platforms. For each ELB, run separate EC2 instance groups to handle the web application for each platform.
- D. Assign multiple ELBs to an EC2 instance or group of EC2 instances running the common components of the web application, one ELB for each platform type. Session stickiness and SSL termination are done at the ELBs.

**Answer:** D

**Explanation:**

One ELB cannot handle different SSL certificates, but since we are using sticky sessions, it must be handled at the ELB level. SSL could be handled on the EC2 instances only with TCP configured ELB.

ELB supports sticky sessions only in HTTP/HTTPS configurations.

The way the Elastic Load Balancer does session stickiness is on an HTTP/HTTPS listener by utilizing an HTTP cookie. If SSL traffic is not terminated on the Elastic Load Balancer and is

terminated on the back-end instance, the Elastic Load Balancer has no visibility into the HTTP headers and therefore cannot set or read any of the HTTP headers.

being passed back and forth. <http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/elb-sticky-sessions.html>

**NEW QUESTION 8**

You're running an application on-premises due to its dependency on non-x86 hardware and want to use AWS for data backup. Your backup application is only able to write to POSIX-compatible blockbased storage. You have 140TB of data and would like to mount it as a single folder on your file server. Users must be able to access portions of this data while the backups are taking place. What backup solution would be most appropriate for this use case?

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 9**

Your website is serving on-demand training videos to your workforce. Videos are uploaded monthly in high resolution MP4 format. Your workforce is distributed globally often on the move and using company-provided tablets that require the HTTP Live Streaming (HLS) protocol to watch a video. Your company has no video transcoding expertise and it required you may need to pay for a consultant. How do you implement the most cost-efficient architecture without compromising high availability and quality of video delivery'?

- A. A video transcoding pipeline running on EC2 using SQS to distribute tasks and Auto Scaling to adjust the number of nodes depending on the length of the queue
- B. EBS volumes to host videos and EBS snapshots to incrementally backup original files after a few day
- C. CloudFront to serve HLS transcoded videos from EC2.
- D. Elastic Transcoder to transcode original high-resolution MP4 videos to HL
- E. EBS volumes to host videos and EBS snapshots to incrementally backup original files after a few day
- F. CloudFront to serve HLS transcoded videos from EC2.
- G. Elastic Transcoder to transcode original high-resolution MP4 videos to HL
- H. S3 to host videos with Lifecycle Management to archive original files to Glacier after a few day
- I. CloudFront to serve HLS transcoded videos from S3.
- J. A video transcoding pipeline running on EC2 using SQS to distribute tasks and Auto Scaling to adjust the number of nodes depending on the length of the queue
- K. S3 to host videos with Lifecycle Management to archive all files to Glacier after a few day
- L. CloudFront to serve HLS transcoded videos from Glacier.

**Answer:** C

**NEW QUESTION 10**

You are designing a connectivity solution between on-premises infrastructure and Amazon VPC. Your server's on-premises will be communicating with your VPC instances. You will be establishing IPsec tunnels over the Internet. You will be using VPN gateways and terminating the IPsec tunnels on AWS-supported customer gateways. Which of the following objectives would you achieve by implementing an IPsec tunnel as outlined above? (Choose four.)

- A. End-to-end protection of data in transit
- B. End-to-end Identity authentication
- C. Data encryption across the Internet
- D. Protection of data in transit over the Internet
- E. Peer identity authentication between VPN gateway and customer gateway
- F. Data integrity protection across the Internet

**Answer:** CDEF

**NEW QUESTION 10**

You are designing an intrusion detection prevention (IDS/IPS) solution for a customer web application in a single VPC. You are considering the options for implementing IDS/IPS protection for traffic coming from the Internet. Which of the following options would you consider? (Choose two.)

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

**Answer:** AD

**Explanation:**

EC2 does not allow promiscuous mode, and you cannot put something in between the ELB and the web server (like a listener or IDP)

**NEW QUESTION 15**

You are designing a photo sharing mobile app the application will store all pictures in a single Amazon S3 bucket. Users will upload pictures from their mobile device directly to Amazon S3 and will be able to view and download their own pictures directly from Amazon S3. You want to configure security to handle potentially millions of users in the most secure manner possible. What should your server-side application do when a new user registers on the photosharing mobile application?

- A. Create a set of long-term credentials using AWS Security Token Service with appropriate permissions. Store these credentials in the mobile app and use them to

access Amazon S3.

B. Record the user's Information in Amazon RDS and create a role in IAM with appropriate permission

C. When the user uses their mobile app create temporary credentials using the AWS Security Token Service 'AssumeRole' function Store these credentials in the mobile app's memory and use them to access Amazon S3 Generate new credentials the next time the user runs the mobile app.

D. Record the user's Information In Amazon DynamoD

E. When the user uses their mobile app create temporary credentials using AWS Security Token Service with appropriate permissions Store these credentials in the mobile app's memory and use them to access Amazon S3 Generate newcredentials the next time the user runs the mobile app.

F. Create IAM use

G. Assign appropriate permissions to the IAM user Generate an access key and secret key for the IAM user, store them in the mobile app and use these credentials to access Amazon S3.

H. Create an IAM use

I. Update the bucket policy with appropriate permissions for the IAM user Generate an access Key and secret Key for the IAM user, store them In the mobile app and use these credentials to access Amazon S3.

**Answer: C**

#### NEW QUESTION 19

Your company is getting ready to do a major public announcement of a social media site on AWS. The website is running on EC2 instances deployed across multiple Availability Zones with a Multi-AZ RDS MySQL Extra Large DB Instance. The site performs a high number of small reads and writes per second and relies on an eventual consistency model. After comprehensive tests you discover that there is read contention on RDS MySQL. Which are the best approaches to meet these requirements? (Choose two.)

A. Deploy ElasticCache in-memory cache running in each availability zone

B. Implement sharding to distribute load to multiple RDS MySQL instances

C. Increase the RDS MySQL Instance size and Implement provisioned IOPS

D. Add an RDS MySQL read replica in each availability zone

**Answer: AD**

#### NEW QUESTION 23

To serve Web traffic for a popular product your chief financial officer and IT director have purchased 10 ml 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 ml 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 Autoscaning group and Launch Configuration with ELB to add up to 10 more on-demand mi large instances when triggered by Cloudwatch shut off c3 2xiarge instances

C. Route traffic to EC2 ml 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 mi .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 26

You are running a news website in the eu-west-1 region that updates every 15 minutes. The website has a world-wide audience it uses an Auto Scaling group behind an Elastic Load Balancer and an Amazon RDS database Static content resides on Amazon S3, and is distributed through Amazon CloudFront. Your Auto Scaling group is set to trigger a scale up event at 60% CPU utilization, you use an Amazon RDS extra large DB instance with 10.000 Provisioned IOPS its CPU utilization is around 80%. While freeable memory is in the 2 GB range.

Web analytics reports show that the average load time of your web pages is around 1 5 to 2 seconds,

but your SEO consultant wants to bring down the average load time to under 0.5 seconds. How would you improve page load times for your users? (Choose three.)

A. Lower the scale up trigger of your Auto Scaling group to 30% so it scales more aggressively.

B. Add an Amazon ElastiCache caching layer to your application for storing sessions and frequent DB queries

C. Configure Amazon CloudFront dynamic content support to enable caching of re-usable content from your site

D. Switch Amazon RDS database to the high memory extra large Instance type

E. Set up a second installation in another region, and use the Amazon Route 53 latency-based routing feature to select the right region.

**Answer: BCE**

#### Explanation:

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

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

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

the innodb\_buffer\_pool\_size (by default set to 75% of physical memory) is one way example of adjusting MySQL memory settings

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

**NEW QUESTION 31**

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

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

**Answer:** D

**NEW QUESTION 32**

You are looking to migrate your Development (Dev) and Test environments to AWS. You have decided to use separate AWS accounts to host each environment. You plan to link each accounts bill to a Master AWS account using Consolidated Billing. To make sure you Keep within budget you would like to implement a way for administrators in the Master account to have access to stop, delete and/or terminate resources in both the Dev and Test accounts. Identify which option will allow you to achieve this goal.

- A. Create IAM users in the Master account with full Admin permission
- B. Create cross-account roles in the Dev and Test accounts that grant the Master account access to the resources in the account by inheriting permissions from the Master account.
- C. Create IAM users and a cross-account role in the Master account that grants full Admin permissions to the Dev and Test accounts.
- D. Create IAM users in the Master account
- E. Create cross-account roles in the Dev and Test accounts that have full Admin permissions and grant the Master account access.
- F. Link the accounts using Consolidated Billing
- G. This will give IAM users in the Master account access to resources in the Dev and Test accounts

**Answer:** C

**NEW QUESTION 37**

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 39**

You deployed your company website using Elastic Beanstalk and you enabled log file rotation to S3. An Elastic Map Reduce job is periodically analyzing the logs on S3 to build a usage dashboard that you share with your CIO.

You recently improved overall performance of the website using Cloud Front for dynamic content delivery and your website as the origin.

After this architectural change, the usage dashboard shows that the traffic on your website dropped by an order of magnitude. How do you fix your usage dashboard'?

- A. Enable Cloud Front to deliver access logs to S3 and use them as input of the Elastic Map Reduce job.
- B. Turn on Cloud Trail and use trail log files on S3 as input of the Elastic Map Reduce job
- C. Change your log collection process to use Cloud Watch ELB metrics as input of the Elastic Map Reduce job
- D. Use Elastic Beanstalk "Rebuild Environment" option to update log delivery to the Elastic Map Reduce job.
- E. Use Elastic Beanstalk "Restart App server(s)" option to update log delivery to the Elastic Map Reduce job.

**Answer:** A

**NEW QUESTION 42**

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

- A. Configure an instance with monitoring software and the elastic network interface (ENI) set to promiscuous mode packet sniffing to see an traffic across the VPC,
- B. Create a second VPC and route all traffic from the primary application VPC through the second VPC where the scalable virtualized IDS/IPS platform resides.

- C. Configure servers running in the VPC using the host-based 'route' commands to send all traffic through the platform to a scalable virtualized IDS/IPS.  
D. Configure each host with an agent that collects all network traffic and sends that traffic to the IDS/IPS platform for inspection.

**Answer:** D

**Explanation:**

- \A. Not possible to set an instance's NIC into promiscuous mode.  
\B. Incorrect... VPC peering connections are not "transitive", i.e. you cannot pass traffic through a VPC peering connection into another VPC, and then have that other VPC send the traffic to some third VPC, or the Internet, or a VPN, or a direct connect circuit. (I would assume AWS does not allow redistribution of routes from one VPC's back-end VRF into another VPC's back-end VRF, unless it is that first VPC's CIDR block? Someone from AWS would have to chime in here, and they're probably not going to tell us.)  
\C. This one is incorrect because adding static routes on an instance won't affect the routing from any point after the packet leaves the instance's NIC. AWS will check the destination IP address in the packet header and forward according to the VPC routing table's routes. You'd need to make routing changes in the VPC route table for that instance's traffic to get sent through another device (e.g. NAT gateway, VPN instance, or security proxy in this case). (You could tunnel/proxy the traffic over through the IPS tier by changing the destination IP address in the IP header of the packet before it left the instance. But choice C did not state anything about doing anything like that. It just said add a static route on the instance, which does not change the destination IP address in the IP header of the packet.)  
\D. Correct, this is the standard approach, and is definitely scalable.

**NEW QUESTION 43**

A web-startup runs its very successful social news application on Amazon EC2 with an Elastic Load Balancer, an Auto-Scaling group of Java/Tomcat application-servers, and DynamoDB as data store. The main web-application best runs on m2 x large instances since it is highly memory- bound Each new deployment requires semi-automated creation and testing of a new AMI for the application servers which takes quite a while and is therefore only done once per week. Recently, a new chat feature has been implemented in nodejs and waits to be integrated in the architecture. First tests show that the new component is CPU bound Because the company has some experience with using Chef, they decided to streamline the deployment process and use AWS Ops Works as an application life cycle tool to simplify management of the application and reduce the deployment cycles.

What configuration in AWS Ops Works is necessary to integrate the new chat module in the most cost-efficient and filexible way?

- A. Create one AWS OpsWorks stack, create one AWS Ops Works layer, create one custom recipe  
B. Create one AWS OpsWorks stack create two AWS Ops Works layers create one custom recipe  
C. Create two AWS OpsWorks stacks create two AWS Ops Works layers create one custom recipe  
D. Create two AWS OpsWorks stacks create two AWS Ops Works layers create two custom recipe

**Answer:** B

**Explanation:**

You only need one stack to contain two layers:

- one layer for the Java/Tomcat instances
- one layer for DynamoDB

You'd only need one custom recipe because the only OpsWorks Lifecycle Event that would be involved in rolling out the new chat feature would be "Deploy". (Or you could implement it in "Setup" if you choose to make including the chat app a new baseline standard for your instances in that layer. But even then, you'd only have one custom recipe because there would be no need to customize the "Deploy" event to install the chat app if you already installed out the chat app in "Setup".) So you'd need a custom recipe for that one lifecycle event. And it would only be used for the "Deploy" lifecycle event on the app layer, not on the DB layer

**NEW QUESTION 47**

Before I delete an EBS volume, what can I do if I want to recreate the volume later?

- A. Create a copy of the EBS volume (not a snapshot)  
B. Store a snapshot of the volume  
C. Download the content to an EC2 instance  
D. Back up the data in to a physical disk

**Answer:** B

**Explanation:**

After you no longer need an Amazon EBS volume, you can delete it. After deletion, its data is gone and the volume can't be attached to any instance. However, before deletion, you can store a snapshot of the volume, which you can use to re-create the volume later.

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

**NEW QUESTION 50**

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 53**

What is the Reduced Redundancy option in Amazon S3?

- A. Less redundancy for a lower cost.  
B. It doesn't exist in Amazon S3, but in Amazon EBS.  
C. It allows you to destroy any copy of your files outside a specific jurisdiction.  
D. It doesn't exist at all

**Answer:** A

**NEW QUESTION 58**

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 63**

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

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

**Answer:** D

**Explanation:**

When you view the block device mapping for your instance, you can see only the EBS volumes, not the instance store volumes. You can use instance metadata to query the complete block device mapping. The base URI for all requests for instance metadata is <http://169.254.169.254/latest/>. <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/block-device-mappingconcepts.html#bdm-instance-metadata>

**NEW QUESTION 67**

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

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

**Answer:** B

**Explanation:**

Q: Will my standby be in the same Region as my primary?

Yes. Your standby is automatically provisioned in a **different Availability Zone of the same Region** as your DB instance primary.

**NEW QUESTION 68**

What does Amazon Elastic Beanstalk provide?

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

**Answer:** B

**NEW QUESTION 69**

What does the AWS Storage Gateway provide?

- A. It allows to integrate on-premises IT environments with Cloud Storage.
- B. A direct encrypted connection to Amazon S3.
- C. It's a backup solution that provides an on-premises Cloud storage.
- D. It provides an encrypted SSL endpoint for backups in the Clou

**Answer:** A

**NEW QUESTION 74**

What is the maximum key length of a tag?

- A. 512 Unicode characters
- B. 64 Unicode characters
- C. 256 Unicode characters
- D. 128 Unicode characters

**Answer:** D

**NEW QUESTION 76**

You must increase storage size in increments of at least \_\_\_\_\_ %

- A. 40
- B. 20

- C. 50
- D. 10

**Answer:** D

**NEW QUESTION 80**

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

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

**Answer:** A

**Explanation:**

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

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

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

**NEW QUESTION 84**

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 87**

What are the Amazon EC2 API tools?

- A. They don't exist
- B. The Amazon EC2 AMI tools, instead, are used to manage permissions.
- C. Command-line tools to the Amazon EC2 web service.
- D. They are a set of graphical tools to manage EC2 instances.
- E. They don't exist
- F. The Amazon API tools are a client interface to Amazon Web Services.

**Answer:** B

**NEW QUESTION 91**

Can Amazon S3 uploads resume on failure or do they need to restart?

- A. Restart from beginning
- B. You can resume them, if you flag the "resume on failure" option before uploading.
- C. Resume on failure
- D. Depends on the file size

**Answer:** C

**NEW QUESTION 95**

Fill in the blanks: \_\_\_\_\_ let you categorize your EC2 resources in different ways, for example, by purpose, owner, or environment.

- A. wildcards
- B. pointers
- C. Tags
- D. special filters

**Answer:** C

**NEW QUESTION 99**

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

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

**Answer:** A

**Explanation:**

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

**Security Groups for Load Balancers in a VPC**

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

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

[Default VPC] If you use the **AWS CLI or API** to create a load balancer in your default VPC, you can't choose an existing security group for your load balancer. Instead, Elastic Load Balancing provides a security group with rules to allow all traffic on the ports specified for the load balancer. Elastic Load Balancing creates only one such security group per AWS account, with a name of the form default\_elb\_id (for example, default\_elb\_fc5fbed3-0405-3b7d-a328-aa290EXAMPLE). Subsequent load balancers that you create in the default VPC also use this security group. Be sure to review the security group rules to ensure that they allow traffic on the listener and health check ports for the new load balancer. When you delete your load balancer, this security group is not deleted automatically.

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

**NEW QUESTION 102**

What is the durability of S3 RRS?

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

**Answer:** A

**Explanation:**

RRS = Reduced Redundancy Storage

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

**NEW QUESTION 106**

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 110**

A/An \_\_\_\_ acts as a firewall that controls the traffic allowed to reach one or more instances.

- A. security group
- B. ACL
- C. IAM
- D. Private IP Addresses

**Answer:** A

**Explanation:**

A security group acts as a virtual firewall that controls the traffic for one or more instances. <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html>

**NEW QUESTION 111**

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

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

**Answer:** D

**Explanation:**

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

**NEW QUESTION 115**

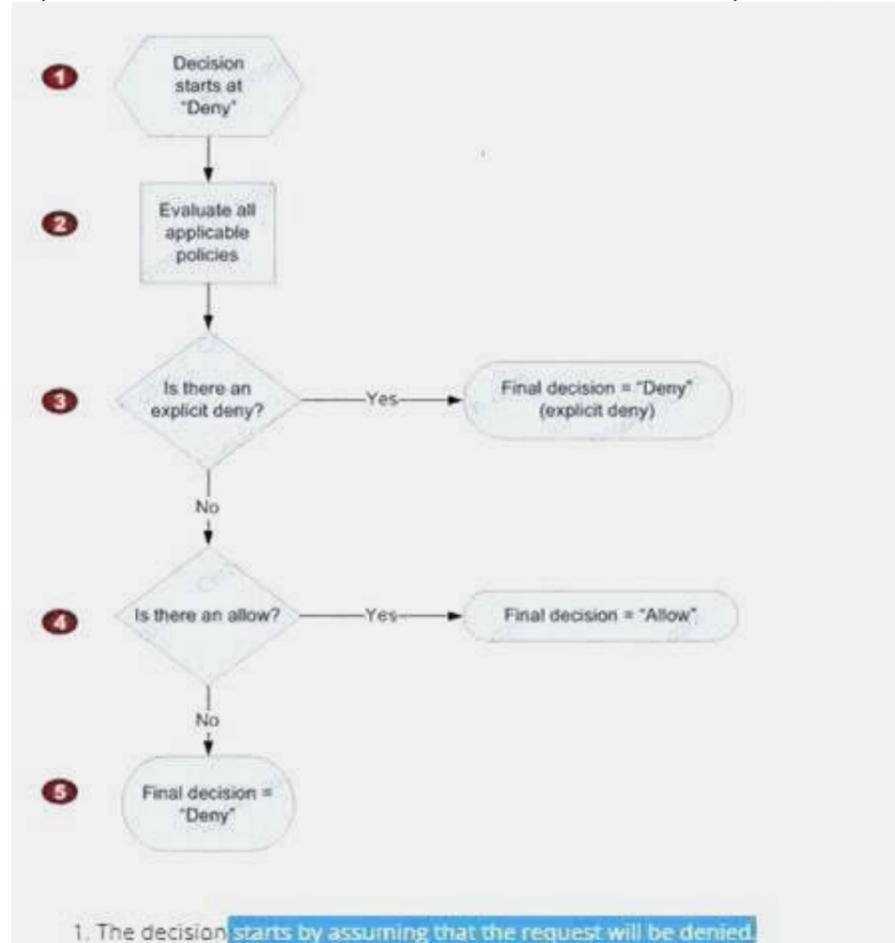
IAM's Policy Evaluation Logic always starts with a default \_\_\_\_ for every request, except for those that use the AWS account's root security credentials b

- A. Permit
- B. Deny
- C. Cancel

**Answer:** B

**Explanation:**

[http://docs.aws.amazon.com/IAM/latest/UserGuide/reference\\_policies\\_evaluation-logic.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_evaluation-logic.html)



**NEW QUESTION 120**

For each DB Instance class, what is the maximum size of associated storage capacity?

- A. 5GB
- B. 6TB
- C. 2TB
- D. 500GB

**Answer:** B

**Explanation:**

"You can now create MySQL, PostgreSQL, and Oracle RDS database instances with up to 6TB of storage and SQL Server RDS database instances with up to 4TB of storage when using the Provisioned IOPS and General Purpose (SSD) storage types. Existing MySQL, PostgreSQL, and Oracle RDS database instances can be scaled to these new database storage limits without any downtime."

**NEW QUESTION 125**

While signing in REST/ Query requests, for additional security, you should transmit your requests using Secure Sockets Layer (SSL) by using \_\_\_\_

- A. HTTP
- B. Internet Protocol Security(IPsec)
- C. TLS (Transport Layer Security)
- D. HTTPS

**Answer:** D

**NEW QUESTION 127**

What happens to the I/O operations while you take a database snapshot?

- A. I/O operations to the database are suspended for a few minutes while the backup is in progress.
- B. I/O operations to the database are sent to a Replica (if available) for a few minutes while the backup is in progress.
- C. I/O operations will be functioning normally

D. I/O operations to the database are suspended for an hour while the backup is in progress

**Answer:** A

**Explanation:**

Creating this DB snapshot on a Single-AZ DB instance results in a brief I/O suspension that typically lasting no more than a few minutes. Multi-AZ DB instances are not affected by this I/O suspension since the backup is taken on the standby.

**NEW QUESTION 132**

When running my DB Instance as a Multi-AZ deployment, can I use the standby for read or write operations?

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

**Answer:** D

**Explanation:**

Q: When running my DB instance as a Multi-AZ deployment, can I use the standby for read or write operations?

**No**, the standby replica cannot serve read requests. Multi-AZ deployments are designed to provide enhanced database availability and durability, rather than read scaling benefits. As such, the feature uses synchronous replication between primary and standby. Our implementation makes sure the primary and the standby are constantly in sync, but precludes using the standby for read or write operations. If you are interested in a read scaling solution, please see the FAQs on Read Replicas.

**NEW QUESTION 135**

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

- A. USD 0.10 per GB
- B. No charge
- C. It is free.
- D. USD 0.02 per GB
- E. USD 0.01 per GB

**Answer:** B

**Explanation:**

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

**NEW QUESTION 137**

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

- A. AWS Access Control Service (ACS)
- B. AWS Identity and Access Management (IAM)
- C. AWS Identity Manager (AIM)
- D. AWS Security Groups

**Answer:** B

**NEW QUESTION 140**

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

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 143**

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

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

**Answer:** B

**Explanation:**

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

**NEW QUESTION 147**

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

- A. FALSE
- B. TRUE

**Answer: B**

**Explanation:**

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

**NEW QUESTION 148**

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

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

**Answer: B**

**Explanation:**

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

**NEW QUESTION 149**

Can I encrypt connections between my application and my DB Instance using SSL?

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

**Answer: B**

**NEW QUESTION 152**

Please select the Amazon EC2 resource which cannot be tagged.

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

**Answer: C**

**Explanation:**

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

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

**NEW QUESTION 154**

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

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

**Answer:** A

**NEW QUESTION 155**

Select the incorrect statement

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

**Answer:** C

**Explanation:**

A private IP address remains associated with the network interface when the instance is stopped and restarted, and is released when the instance is terminated.

**NEW QUESTION 160**

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

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

**Answer:** A

**Explanation:**

Monitoring Volumes with CloudWatch

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

Basic

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

Detailed

Provisioned IOPS SSD (io1) volumes automatically send one-minute metrics to CloudWatch.

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

### Monitoring Volumes with CloudWatch

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

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

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

#### NEW QUESTION 161

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 sense.

**Answer: B**

#### NEW QUESTION 165

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

- A. FALSE
- B. TRUE

**Answer: B**

#### Explanation:

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

#### NEW QUESTION 167

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

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

**Answer: B**

#### Explanation:

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

#### NEW QUESTION 171

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

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

**Answer: A**

#### Explanation:

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

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

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

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

### Default Security Groups

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

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

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

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

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

#### NEW QUESTION 175

What does Amazon Route53 provide?

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

**Answer: C**

#### Explanation:

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

#### NEW QUESTION 180

What does Amazon ElastiCache provide?

- A. A service by this name doesn't exist
- B. Perhaps you mean Amazon CloudCache.
- C. A virtual server with a huge amount of memory.
- D. A managed In-memory cache service.
- E. An Amazon EC2 instance with the Memcached software already pre-installed

**Answer: C**

#### NEW QUESTION 181

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. A firewall for inbound traffic, built-in around every Amazon EC2 instance

**Answer: D**

#### Explanation:

A security group acts as a virtual firewall for your instance to control inbound and outbound traffic.  
[http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC\\_SecurityGroups.html](http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_SecurityGroups.html)

#### NEW QUESTION 184

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

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

**Answer: AC**

#### NEW QUESTION 189

Amazon EC2 has no Amazon Resource Names (ARNs) because you can't specify a particular Amazon EC2 resource in an IAM policy.

- A. TRUE
- B. FALSE

**Answer:** B

**Explanation:**

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

**NEW QUESTION 192**

Is the encryption of connections between my application and my DB Instance using SSL for the MySQL server engines available?

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

**Answer:** A

**Explanation:**

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

**Q: Can I encrypt connections between my application and my DB Instance using SSL?**  
Yes, this option is currently supported for the MySQL, MariaDB, SQL Server, PostgreSQL, and Oracle engines.  
**Amazon RDS generates an SSL certificate for each DB Instance.** Once an encrypted connection is established, data transferred between the DB Instance and your application will be encrypted during transfer.

**NEW QUESTION 195**

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

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

**Answer:** A

**Explanation:**

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

**Amazon EC2 Instance Store**  
Many instances can access storage from disks that are physically attached to the host computer. This disk storage is referred to as *instance store*. Instance store provides temporary block-level storage for instances. **The data on an instance store volume persists only during the life of the associated instance. If you stop or terminate an instance, any data on instance store volumes is lost.** For more information, see Amazon EC2 Instance Store.

**NEW QUESTION 198**

Security groups act like a firewall at the instance level, whereas \_\_\_\_\_ are an additional layer of security that act at the subnet level.

- A. DB Security Groups
- B. VPC Security Groups
- C. network ACLs

**Answer:** C

**NEW QUESTION 203**

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

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

**Answer:** A

**Explanation:**

Q: My Amazon RDS for MySQL Read Replica appears "stuck" after a Multi-AZ failover and is unable to obtain or apply updates from the source DB Instance. What do I do? ... To resolve the current issue, you will need to delete the Read Replica and create a new one to replace it. " <https://aws.amazon.com/rds/faqs/>

**NEW QUESTION 207**

If your DB instance runs out of storage space or file system resources, its status will change to \_\_\_\_\_ and your DB Instance will no longer be available.

- A. storage-overflow
- B. storage-full
- C. storage-exceed
- D. storage-overage

**Answer:** B

**Explanation:**

<https://aws.amazon.com/ko/premiumsupport/knowledge-center/rds-out-of-storage/>

**Short Description**

When an RDS DB instance reaches the **STORAGE\_FULL** state, there is **not enough space available** for performing basic operations, eventually preventing you from restarting or making connections to the instance.

**NEW QUESTION 211**

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

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

**Answer:** A

**NEW QUESTION 216**

Which Amazon storage do you think is the best for my database-style applications that frequently encounter many random reads and writes across the dataset?

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

**Answer:** D

**Explanation:**

"Amazon EBS is particularly helpful for database-style applications that frequently encounter many random reads and writes across the data set."  
<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonEBS.html>

**NEW QUESTION 219**

In the context of MySQL, version numbers are organized as MySQL version = X.Y.Z. What does X denote here?

- A. release level
- B. minor version
- C. version number
- D. major version

**Answer:** D

**Explanation:**

**MySQL on Amazon RDS Versions**

For MySQL, version numbers are organized as version = X.Y.Z. In Amazon RDS terminology, **X.Y denotes the major version**, and **Z is the minor version number**. For Amazon RDS implementations, a version change is considered major if the major version number changes—for example, going from version 5.6 to 5.7. A version change is considered minor if only the minor version number changes—for example, going from version 5.6.22 to 5.6.23.

Amazon RDS currently supports MySQL major versions 5.5, 5.6, and 5.7. MySQL minor version support varies by AWS Region. Use the following table to see what MySQL minor versions are supported in each AWS Region.

**NEW QUESTION 220**

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

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

**Answer:** C

**NEW QUESTION 224**

It is advised that you watch the Amazon CloudWatch " \_\_\_\_\_ " metric (available via the AWS Management Console or Amazon Cloud Watch APIs) carefully and recreate the Read Replica should it fall behind due to replication errors.

- A. Write Lag
- B. Read Replica
- C. Replica Lag
- D. Single Replica

**Answer:** C

**Explanation:**

The amount of time a Read Replica DB instance lags behind the source DB instance. Applies to MySQL, MariaDB, and PostgreSQL Read Replicas.  
<http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/rds-metricscollected.html>

<b>ReplicaLag</b>	The amount of time a Read Replica DB instance lags behind the source DB instance. Applies to MySQL, MariaDB, and PostgreSQL Read Replicas.  Units: Seconds
-------------------	--

**NEW QUESTION 226**

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

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

**Answer: C**

**Explanation:**

Private MFA does not apply here.

Q. What is AWS MFA?

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

**NEW QUESTION 227**

You can use \_\_\_\_\_ and \_\_\_\_\_ to help secure the instances in your VPC,

- A. security groups and multi-factor authentication
- B. security groups and 2-Factor authentication
- C. security groups and biometric authentication
- D. security groups and network ACLs

**Answer: D**

**NEW QUESTION 232**

Is there a limit to the number of groups you can have?

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

**Answer: D**

**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. [http://docs.aws.amazon.com/IAM/latest/UserGuide/reference\\_iam-limits.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_iam-limits.html)

**NEW QUESTION 237**

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 240**

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

- A. 30 minutes
- B. 1 hour
- C. 12 hours
- D. 10 minutes

**Answer: B**

**NEW QUESTION 244**

Does Dynamic DB support in-place atomic updates?

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

**Answer:** C

**Explanation:**

Q: Does DynamoDB support in-place atomic updates?

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

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

**NEW QUESTION 247**

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

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

**Answer:** C

**Explanation:**

Amazon EC2 uses SSH keys, Windows passwords, and security groups to control who has access to the operating system of specific Amazon EC2 instances. There's no method in the IAM system to allow or deny access to the operating system of a specific instance.

[http://docs.aws.amazon.com/IAM/latest/UserGuide/IAM\\_UseCases.html](http://docs.aws.amazon.com/IAM/latest/UserGuide/IAM_UseCases.html)

**NEW QUESTION 252**

What's an ECU?

- A. Extended Cluster User.
- B. None of these.
- C. Elastic Computer Usage.
- D. Elastic Compute Uni

**Answer:** B

**Explanation:**

The EC2 Compute Unit (ECU) provides the relative measure of the integer processing power of an Amazon EC2 instance.

<https://aws.amazon.com/ec2/faqs/>

**NEW QUESTION 257**

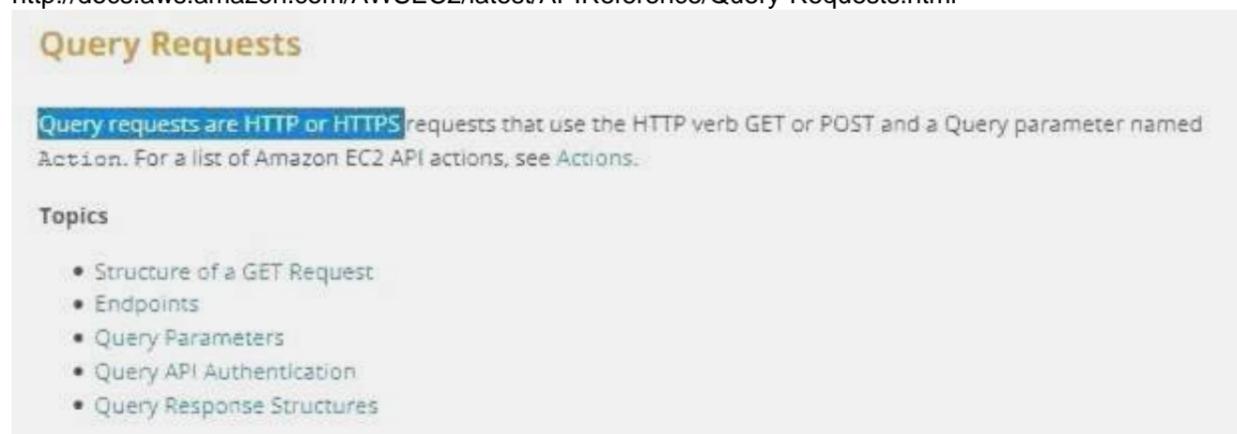
REST or Query requests are HTTP or HTTPS requests that use an HTTP verb (such as GET or POST) and a parameter named Action or Operation that specifies the API you are calling.

- A. FALSE
- B. TRUE

**Answer:** B

**Explanation:**

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



**NEW QUESTION 259**

Within the IAM service a GROUP is regarded as a:

- A. A collection of AWS accounts
- B. It's the group of EC2 machines that gain the permissions specified in the GROUP.
- C. There's no GROUP in IAM, but only USERS and RESOURCES.
- D. A collection of user

**Answer:** D

**Explanation:**

Use groups to assign permissions to IAM users

Instead of defining permissions for individual IAM users, it's usually more convenient to create groups that relate to job functions (administrators, developers, accounting, etc.), define the relevant permissions for each group, and then assign IAM users to those groups. All the users in an IAM group inherit the permissions assigned to the group. That way, you can make changes for everyone in a group in just one place. As people move around in your company, you can simply change what IAM group their IAM user belongs to.

<http://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html#use-groups-forpermissions>

**NEW QUESTION 264**

A \_\_\_\_\_ is the concept of allowing (or disallowing) an entity such as a user, group, or role some type of access to one or more resources.

- A. user
- B. AWS Account
- C. resource
- D. permission

**Answer: D**

**Explanation:**

A permission is the concept of allowing (or disallowing) an entity such as a user, group, or role some type of access to one or more resources.

**NEW QUESTION 267**

After an Amazon VPC instance is launched, can I change the VPC security groups it belongs to?

- A. N
- B. You cannot.
- C. Ye
- D. You can.
- E. Only if you are the root user
- F. Only if the tag "VPC\_Change\_Group" is true

**Answer: B**

**Explanation:**

Security groups are associated with network interfaces. After you launch an instance, you can change the security groups associated with the instance, which changes the security groups associated with the primary network interface (eth0).

**NEW QUESTION 268**

Do the system resources on the Micro instance meet the recommended configuration for Oracle?

- A. Yes completely
- B. Yes but only for certain situations
- C. Not in any circumstance

**Answer: C**

**Explanation:**

We recommend that you use db.t1.micro instances with Oracle to test setup and connectivity only; the system resources for a db.t1.micro instance do not meet the recommended configuration for Oracle. No Oracle options are supported on a db.t1.micro instance.

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.DBInstanceClass.html#Concepts.DBInstanceClasses.Previous>

**NEW QUESTION 271**

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

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

**Answer: C**

**NEW QUESTION 272**

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

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

**Answer: A**

**Explanation:**

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

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

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

**NEW QUESTION 276**

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 277**

What is the charge for the data transfer incurred in replicating data between your primary and standby?

- A. Same as the standard data transfer charge
- B. Double the standard data transfer charge
- C. No charge
- D. It is free
- E. Half of the standard data transfer charge

**Answer:** C

**Explanation:**

Q: How much do Read Replicas cost? When does billing begin and end?

A Read Replica is billed as a standard DB Instance and at the same rates. Click here for more information on DB Instance billing visit this FAQ. Just like a standard DB Instance, the rate per "DB Instance hour" for a Read Replica is determined by the DB Instance class of the Read Replica –please see Amazon RDS detail page for up-to-date pricing. You are not charged for the data transfer incurred in replicating data between your source DB Instance and Read Replica. Billing for a Read Replica begins as soon as the Read Replica has been successfully created (i.e. when status is listed as "active"). The Read Replica will continue being billed at standard Amazon RDS DB Instance hour rates until you issue a command to delete it.

**NEW QUESTION 278**

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 281**

What is the name of licensing model in which I can use your existing Oracle Database licenses to run Oracle deployments on Amazon RDS?

- A. Bring Your Own License
- B. Role Bases License
- C. Enterprise License
- D. License Included

**Answer:** A

**Explanation:**

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

**NEW QUESTION 282**

Regarding the attaching of ENI to an instance, what does 'warm attach' refer to?

- A. Attaching an ENI to an instance when it is stopped.
- B. This question doesn't make sense.
- C. Attaching an ENI to an instance when it is running
- D. Attaching an ENI to an instance during the launch process

**Answer:** A

**Explanation:**

Best Practices for Configuring Elastic Network Interfaces

You can attach an elastic network interface to an instance when it's running (hot attach), when it's stopped (warm attach), or when the instance is being launched (cold attach). <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html#best-practices-forconfiguring-network-interfaces>

**NEW QUESTION 285**

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

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

C. You will be charged for the lowest storage capacity you have used

**Answer:** B

**Explanation:**

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

**NEW QUESTION 290**

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 291**

You are using an m1.small EC2 Instance with one 300 GB EBS volume to host a relational database. You determined that write throughput to the database needs to be increased. Which of the following approaches can help achieve this? (Choose two.)

- A. Use an array of EBS volumes.
- B. Enable Multi-AZ mode.
- C. Place the instance in an Auto Scaling Groups
- D. Add an EBS volume and place into RAID 5.
- E. Increase the size of the EC2 Instance.
- F. Put the database behind an Elastic Load Balance

**Answer:** AE

**NEW QUESTION 294**

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 297**

Which of the following are characteristics of Amazon VPC subnets? (Choose two.)

- A. Each subnet spans at least 2 Availability Zones to provide a high-availability environment.
- B. Each subnet maps to a single Availability Zone.
- C. CIDR block mask of /25 is the smallest range supported.
- D. By default, all subnets can route between each other, whether they are private or public.
- E. Instances in a private subnet can communicate with the Internet only if they have an Elastic I

**Answer:** BD

**Explanation:**

Even though we know the right Answers it is sometimes good to know why the other Answers are wrong.

- A. Is wrong because a subnet maps to a single AZ.
- C. Is wrong because /28 is the smallest subnet, amazon takes first four and last addresses per subnet.
- E. Is wrong because a private subnet needs a NAT appliance.

**NEW QUESTION 302**

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

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

**Answer:** C

**Explanation:**

Objects stored in a region never leave the region unless you explicitly transfer them to another region. For example, objects stored in the EU (Ireland) region never

leave it. <http://docs.aws.amazon.com/AmazonS3/latest/dev/Introduction.html#Regions>

**NEW QUESTION 304**

In AWS, which security aspects are the customer's responsibility? (Choose four.)

- A. Security Group and ACL (Access Control List) settings
- B. Decommissioning storage devices
- C. Patch management on the EC2 instance's operating system
- D. Life-cycle management of IAM credentials
- E. Controlling physical access to compute resources
- F. Encryption of EBS (Elastic Block Storage) volumes

**Answer:** ACDF

**Explanation:**

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

**NEW QUESTION 307**

You have a web application running on six Amazon EC2 instances, consuming about 45% of resources on each instance. You are using auto-scaling to make sure that six instances are running at all times. The number of requests this application processes is consistent and does not experience spikes. The application is critical to your business and you want high availability at all times. You want the load to be distributed evenly between all instances. You also want to use the same Amazon Machine Image (AMI) for all instances. Which of the following architectural choices should you make?

- A. Deploy 6 EC2 instances in one availability zone and use Amazon Elastic Load Balancer.
- B. Deploy 3 EC2 instances in one region and 3 in another region and use Amazon Elastic Load Balancer.
- C. Deploy 3 EC2 instances in one availability zone and 3 in another availability zone and use Amazon Elastic Load Balancer.
- D. Deploy 2 EC2 instances in three regions and use Amazon Elastic Load Balancer.

**Answer:** C

**Explanation:**

A load balancer accepts incoming traffic from clients and routes requests to its registered EC2 instances in one or more Availability Zones.

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/how-elb-works.html> Updated Security Whitepaper link:

<https://d0.awsstatic.com/whitepapers/aws-security-whitepaper.pdf> References:

**NEW QUESTION 312**

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 316**

An instance is launched into a VPC subnet with the network ACL configured to allow all inbound traffic and deny all outbound traffic. The instance's security group is configured to allow SSH from any IP address and deny all outbound traffic. What changes need to be made to allow SSH access to the instance?

- A. The outbound security group needs to be modified to allow outbound traffic.
- B. The outbound network ACL needs to be modified to allow outbound traffic.
- C. Nothing, it can be accessed from any IP address using SSH.
- D. Both the outbound security group and outbound network ACL need to be modified to allow outbound traffic.

**Answer:** B

**Explanation:**

Need to open TCP Port 1024-65535 at Outbound Rules

"Allows outbound responses to the remote computer. Network ACLs are stateless, therefore this rule is required to allow response traffic for inbound requests."

[http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC\\_ACLs.html](http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_ACLs.html)

**NEW QUESTION 318**

What are characteristics of Amazon S3? (Choose two.)

- A. S3 allows you to store objects of virtually unlimited size.
- B. S3 offers Provisioned IOPS.
- C. S3 allows you to store unlimited amounts of data.
- D. S3 should be used to host a relational database.
- E. Objects are directly accessible via a UR

**Answer:** CE

**NEW QUESTION 320**

Per the AWS Acceptable Use Policy, penetration testing of EC2 instances:

- A. May be performed by AWS, and will be performed by AWS upon customer request.
- B. May be performed by AWS, and is periodically performed by AWS.
- C. Are expressly prohibited under all circumstances.
- D. May be performed by the customer on their own instances with prior authorization from AWS.
- E. May be performed by the customer on their own instances, only if performed from EC2 instances

**Answer:** D

**Explanation:**

Our Acceptable Use Policy describes permitted and prohibited behavior on AWS and includes descriptions of prohibited security violations and network abuse. However, because penetration testing and other simulated events are frequently indistinguishable from these activities, we have established a policy for customers to request permission to conduct penetration tests and vulnerability scans to or originating from the AWS environment. <http://aws.amazon.com/security/penetration-testing/>

**NEW QUESTION 324**

You are working with a customer who has 10 TB of archival data that they want to migrate to Amazon Glacier. The customer has a 1-Mbps connection to the Internet. Which service or feature provides the fastest method of getting the data into Amazon Glacier?

- A. Amazon Glacier multipart upload
- B. AWS Storage Gateway
- C. VM Import/Export
- D. AWS Import/Export

**Answer:** D

**Explanation:**

You can only perform an Amazon Glacier import from devices of 4 TB in size or smaller.

[https://docs.aws.amazon.com/es\\_es/AWSImportExport/latest/DG/createGlacierimportjobs.html](https://docs.aws.amazon.com/es_es/AWSImportExport/latest/DG/createGlacierimportjobs.html) <http://docs.aws.amazon.com/amazonglacier/latest/dev/uploading-archive-mpu.html>

**NEW QUESTION 328**

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 333**

Which procedure for backing up a relational database on EC2 that is using a set of RAIDed EBS volumes for storage minimizes the time during which the database cannot be written to and results in a consistent backup?

- A. 1. Detach EBS volumes, 2. Start EBS snapshot of volumes, 3. Re-attach EBS volumes
- B. 1. Stop the EC2 Instance
- C. 2. Snapshot the EBS volumes
- D. 1. Suspend disk I/O, 2. Create an image of the EC2 Instance, 3. Resume disk I/O
- E. 1. Suspend disk I/O, 2. Start EBS snapshot of volumes, 3. Resume disk I/O
- F. 1. Suspend disk I/O, 2. Start EBS snapshot of volumes, 3. Wait for snapshots to complete, 4. Resume disk I/O

**Answer:** B

**Explanation:**

<https://aws.amazon.com/cn/premiumsupport/knowledge-center/snapshot-ebs-raid-array/>

To create an "application-consistent" snapshot of your RAID array, stop applications from writing to the RAID array, and flush all caches to disk. Then ensure that the associated EC2 instance is no longer writing to the RAID array by taking steps such as freezing the file system, unmounting the RAID array, or "shutting down the associated EC2 instance". After completing the steps to halt all I/O, take a snapshot of each EBS volume.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebsdetaching-volume.html> You can detach an Amazon EBS volume from an instance explicitly or by terminating the instance. However, if the instance is running, you must first unmount the volume from the instance."

**NEW QUESTION 337**

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

- A. Virtual Private Network connectio
- B. AWS Directory Services, and ClassicLink
- C. Virtual Private Network connectio

- D. AWS Directory Services, and Amazon Workspaces
- E. AWS Directory Service, Amazon Workspaces, and AWS Identity and Access Management
- F. Amazon Elastic Compute Cloud, and AWS Identity and Access Management

**Answer:** B

**Explanation:**

To enable integration, you need to ensure that your domain is reachable via an Amazon Virtual Private Cloud VPC (this could mean that Active Directory domain controllers for your domain are running on Amazon EC2 instances, or that they are reachable via a VPN connection and are located in your on-premises network).

**NEW QUESTION 341**

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

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

**Answer:** AC

**Explanation:**

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

**NEW QUESTION 344**

You run an ad-supported photo sharing website using S3 to serve photos to visitors of your site. At some point you find out that other sites have been linking to the photos on your site, causing loss to your business. What is an effective method to mitigate this?

- A. Remove public read access and use signed URLs with expiry dates.
- B. Use CloudFront distributions for static content.
- C. Block the IPs of the offending websites in Security Groups.
- D. Store photos on an EBS volume of the web server

**Answer:** A

**Explanation:**

A signed URL includes additional information, for example, an expiration date and time, that gives you more control over access to your content.

**NEW QUESTION 346**

When an EC2 instance that is backed by an S3-based AMI is terminated, what happens to the data on the root volume?

- A. Data is automatically saved as an EBS snapshot.
- B. Data is automatically saved as an EBS volume.
- C. Data is unavailable until the instance is restarted.
- D. Data is automatically deleted

**Answer:** D

**Explanation:**

Using the legacy S3 based AMIs, either of the above terminates the instance and you lose all local and ephemeral storage (boot disk and /mnt) forever. Hope you remembered to save the important stuff elsewhere.

**NEW QUESTION 349**

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

- A. Multiple Availability Zones
- B. AWS Direct Connect
- C. EC2 Dedicated Instances
- D. Placement Groups
- E. VPC private subnets

**Answer:** D

**Explanation:**

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

**NEW QUESTION 350**

You have an environment that consists of a public subnet using Amazon VPC and 3 instances that are running in this subnet. These three instances can successfully communicate with other hosts on the Internet. You launch a fourth instance in the same subnet, using the same AMI and security group configuration

you used for the others, but find that this instance cannot be accessed from the Internet. What should you do to enable Internet access?

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

**Answer: B**

**Explanation:**

You launched your instance into a public subnet - a subnet that has a route to an Internet gateway. However, the instance in your subnet also needs a public IP address to be able to communicate with the Internet. By default, an instance in a nondefault VPC is not assigned a public IP address. In this step, you'll allocate an Elastic IP address to your account, and then associate it with your instance.

**NEW QUESTION 354**

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 359**

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 362**

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 366**

A company is deploying a two-tier, highly available web application to AWS. Which service provides durable storage for static content while utilizing lower Overall CPU resources for the web tier?

- A. Amazon EBS volume
- B. Amazon S3
- C. Amazon EC2 instance store
- D. Amazon RDS instance

**Answer: B**

**NEW QUESTION 370**

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

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

**Answer:** D

**NEW QUESTION 374**

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

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

**Answer:** C

**NEW QUESTION 376**

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

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

**Answer:** D

**Explanation:**

Disabling Source/Destination Checks.

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

**NEW QUESTION 377**

Which of the following are use cases for Amazon DynamoDB? (Choose three)

- A. Storing BLOB data.
- B. Managing web sessions.
- C. Storing JSON documents.
- D. Storing metadata for Amazon S3 objects.
- E. Running relational joins and complex updates.
- F. Storing large amounts of infrequently accessed data

**Answer:** BCD

**Explanation:**

Ideal Usage Patterns

- Amazon DynamoDB is ideal for existing or new applications that need a flexible NoSQL database with low read and write latencies, and the ability to scale storage and throughput up or down as needed without code changes or downtime.
- Use cases require a highly available and scalable database because downtime or performance degradation has an immediate negative impact on an organization's business. For e.g. mobile apps, gaming, digital ad serving, live voting and audience interaction for live events, sensor networks, log ingestion, access control for web-based content, metadata storage for Amazon S3 objects, ecommerce shopping carts, and web session management

**NEW QUESTION 379**

A customer implemented AWS Storage Gateway with a gateway-cached volume at their main office. An event takes the link between the main and branch office offline. Which methods will enable the branch office to access their data? (Choose three.)

- A. Use a HTTPS GET to the Amazon S3 bucket where the files are located.
- B. Restore by implementing a lifecycle policy on the Amazon S3 bucket.
- C. Make an Amazon Glacier Restore API call to load the files into another Amazon S3 bucket within four to six hours.
- D. Launch a new AWS Storage Gateway instance AMI in Amazon EC2, and restore from a gateway snapshot.
- E. Create an Amazon EBS volume from a gateway snapshot, and mount it to an Amazon EC2 instance.
- F. Launch an AWS Storage Gateway virtual iSCSI device at the branch office, and restore from a gateway snapshot.

**Answer:** DEF

**Explanation:**

A is certainly not right, because files persisted by Storage Gateway to S3 are not visible, let alone be accessible.

<https://forums.aws.amazon.com/thread.jspa?threadID=109748>

B is invalid option because you cannot apply Lifecycle Policies because AWS Storage Gateway does not give you that option. Cached Volumes are never stored to Glacier and hence "C" is not valid.

**NEW QUESTION 382**

The new DB Instance that is created when you promote a Read Replica retains the backup window period.

- A. TRUE
- B. FALSE

**Answer:** A

**Explanation:**

"The new DB instance that is created when you promote a Read Replica retains the backup retention period, backup window period, and parameter group of the former Read Replica source." [http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER\\_ReadRepl.html](http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html)

**NEW QUESTION 387**

With which AWS orchestration service can you implement Chef recipes?

- A. CloudFormation
- B. Elastic Beanstalk
- C. Opsworks
- D. Lambda

**Answer:** C

**NEW QUESTION 389**

You work for a construction company that has their production environment in AWS. The production environment consists of 3 identical web servers that are launched from a standard Amazon linux AMI using Auto Scaling. The web servers are launched in to the same public subnet and belong to the same security group. They also sit behind the same ELB. You decide to do some test and dev and you launch a 4th EC2 instance in to the same subnet and same security group. Annoyingly your 4th instance does not appear to have internet connectivity. What could be the cause of this?

- A. You need to update your routing table so as to provide a route out for this instance.
- B. Assign an elastic IP address to the fourth instance.
- C. You have not configured a NAT in the public subnet.
- D. You have not configured a routable IP address in the host OS of the fourth instance

**Answer:** C

**NEW QUESTION 390**

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

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

**Answer:** B

**NEW QUESTION 393**

You work for a famous bakery who are deploying a hybrid cloud approach. Their legacy IBM AS400 servers will remain on premise within their own datacenter however they will need to be able to communicate to the AWS environment over a site to site VPN connection. What do you need to do to establish the VPN connection?

- A. Connect to the environment using AWS Direct Connect.
- B. Assign a public IP address to your Amazon VPC Gateway.
- C. Create a dedicated NAT and deploy this to the public subnet.
- D. Update your route table to add a route for the NAT to 0.0.0.0/0.

**Answer:** B

**NEW QUESTION 398**

You are a systems administrator and you need to monitor the health of your production environment. You decide to do this using Cloud Watch, however you notice that you cannot see the health of every important metric in the default dash board. Which of the following metrics do you need to design a custom cloud watch metric for, when monitoring the health of your EC2 instances?

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

**Answer:** B

**NEW QUESTION 402**

You are hosting a MySQL database on the root volume of an EC2 instance. The database is using a large amount of IOPs and you need to increase the IOPs available to it. What should you do?

- A. Migrate the database to an S3 bucket.
- B. Migrate the database to Glacier.
- C. Add 4 additional EBS SSD volumes and create a RAID 10 using these volumes.
- D. Use Cloud Front to cache the databas

**Answer: C**

**NEW QUESTION 404**

You run a website which hosts videos and you have two types of members, premium fee paying members and free members. All videos uploaded by both your premium members and free members are processed by a fileet of EC2 instances which will poll SQS as videos are uploaded. However you need to ensure that your premium fee paying members videos have a higher priority than your free members. How do you design SQS?

- A. SQS allows you to set priorities on individual items within the queue, so simply set the fee paying members at a higher priority than your free members.
- B. Create two SQS queues, one for premium members and one for free member
- C. Program your EC2 fileet to poll the premium queue first and if empty, to then poll your free members SQS queue.
- D. SQS would not be suitable for this scenari
- E. It would be much better to use SNS to encode the videos.

**Answer: B**

**NEW QUESTION 405**

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 407**

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

- A. True
- B. False

**Answer: A**

**NEW QUESTION 412**

In order to enable encryption at rest using EC2 and Elastic Block Store you need to

- A. Configure encryption when creating the EBS volume
- B. Configure encryption using the appropriate Operating Systems file system
- C. Configure encryption using X.509 certificates
- D. Mount the EBS volume in to S3 and then encrypt the bucket using a bucket polic

**Answer: A**

**NEW QUESTION 413**

Amazon S3 provides;

- A. Unlimited File Size for Objects
- B. Unlimited Storage
- C. A great place to run a No SQL database from
- D. The ability to act as a web server for dynamic content (i.
- E. can query a database)

**Answer: B**

**NEW QUESTION 416**

Amazon S3 buckets in all other regions (other than US Standard) provide read-after-write consistency for PUTS of new objects.

- A. True
- B. False

**Answer: A**

**NEW QUESTION 421**

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

- A. True
- B. False

**Answer:** A

**NEW QUESTION 426**

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

- A. True
- B. False

**Answer:** B

**Explanation:**

### Placement Groups

A *placement group* is a logical grouping of instances within a **single Availability Zone**. Placement groups are recommended for applications that benefit from low network latency, high network throughput, or both. To provide the lowest latency, and the highest packet-per-second network performance for your placement group, choose an instance type that supports enhanced networking. For more information, see *Enhanced Networking*.

**NEW QUESTION 429**

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

- A. True
- B. False

**Answer:** B

**NEW QUESTION 434**

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:**

### Placement Groups

A *placement group* is a logical grouping of instances within a single Availability Zone. Placement groups are recommended for applications that benefit from **low network latency**, high network throughput, or both. To provide the lowest latency, and the highest packet-per-second network performance for your placement group, choose an instance type that supports enhanced networking. For more information, see *Enhanced Networking*.

**NEW QUESTION 438**

You are appointed as your company's Chief Security Officer and you want to be able to track all changes made to your AWS environment, by all users and at all times, in all regions. What AWS service should you use to achieve this?

- A. CloudAudit
- B. CloudWatch
- C. CloudTrail
- D. CloudDetective

**Answer:** C

**Explanation:**

## How do I know which user made a particular change to my AWS infrastructure?

### Issue

I want to track which users are making changes to my AWS resources and infrastructure. How do I do this?

### Resolution

Although AWS doesn't track this information by default, you can enable AWS CloudTrail for your resources, which will create logs of API calls made on your account and deliver them to an S3 bucket you specify. This will allow you to track changes to your resources, and see which user made the changes.

For more information about setting up CloudTrail, see [Getting Started with CloudTrail](#).

### Keywords

CloudTrail, API, log

### NEW QUESTION 441

What are the different types of virtualization available on EC2?

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

**Answer: B**

### NEW QUESTION 444

Which of the following services allows you root access (i.e. you can login using SSH)?

- A. Elastic Load Balancer
- B. Elastic Map Reduce
- C. ElastiCache
- D. RDS

**Answer: B**

### Explanation:

When you use SSH with AWS, you are connecting to an EC2 instance, which is a virtual server running in the cloud. When working with Amazon EMR, the most common use of SSH is to connect to the EC2 instance that is acting as the master node of the cluster.

### NEW QUESTION 447

You are a solutions architect working for a large digital media company. Your company is migrating their production estate to AWS and you are in the process of setting up access to the AWS console using Identity Access Management (IAM). You have created 5 users for your system administrators. What further steps do you need to take to enable your system administrators to get access to the AWS console?

- A. Generate an Access Key ID & Secret Access Key, and give these to your system administrators.
- B. Enable multi-factor authentication on their accounts and define a password policy.
- C. Generate a password for each user created and give these passwords to your system administrators.
- D. Give the system administrators the secret access key and access key id, and tell them to use these credentials to log in to the AWS console.

**Answer: C**

### NEW QUESTION 448

In Identity and Access Management, when you first create a new user, certain security credentials are automatically generated. Which of the below are valid security credentials?

- A. Access Key ID, Authorized Key
- B. Private Key, Secret Access Key
- C. Private Key, Authorized Key
- D. Access Key ID, Secret Access Key

**Answer: D**

### NEW QUESTION 451

What are the valid methodologies for encrypting data on S3?

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

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

**Answer:** A

**NEW QUESTION 453**

After an Amazon EC2-VPC instance is launched, can I change the VPC security groups it belongs to?

- A. No
- B. Yes
- C. Only if you are the root user
- D. Only if the tag "VPC\_Change\_Group" is true

**Answer:** B

**NEW QUESTION 456**

What is the default VPC security group limit?

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

**Answer:** A

**NEW QUESTION 457**

What does ec2-create-group do with respect to the Amazon EC2 security groups?

- A. Creates a new rule inside the security group.
- B. Creates a new security group for use with your account.
- C. Creates a new group inside the security group.
- D. Groups the user created security groups in to a new group for easy acces

**Answer:** B

**NEW QUESTION 460**

By default, what happens to ENIs that are automatically created and attached to EC2 instances when the attached instance terminates?

- A. Remain as is
- B. Terminate
- C. Hibernate
- D. Pause

**Answer:** B

**Explanation:**

By default, elastic network interfaces that are automatically created and attached to instances using the console are set to terminate when the instance terminates. However, network interfaces created using the command line interface aren't set to terminate when the instance terminates.

Source: [http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html#change\\_term\\_behavior](http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html#change_term_behavior)

**NEW QUESTION 462**

In a management network scenario, which interface on the instance handles public-facing traffic?

- A. Primary network interface
- B. Subnet interface
- C. Secondary network interface

**Answer:** C

**NEW QUESTION 463**

While creating an EC2 snapshot using the API, which Action should I be using?

- A. MakeSnapShot
- B. FreshSnapshot
- C. DeploySnapshot
- D. CreateSnapshot

**Answer:** D

**NEW QUESTION 466**

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

- A. True

B. False

**Answer:** A

**NEW QUESTION 468**

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 471**

What is the maximum write throughput I can provision per table for a single DynamoDB table?

- A. 5,000 us east, 1,000 all other regions
- B. 100,000 us east, 10, 000 all other regions
- C. Designed to scale without limits, but if you go beyond 40,000 us east/10,000 all other regions you have to contact AWS first.
- D. There is no limit

**Answer:** C

**NEW QUESTION 475**

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

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

**Answer:** C

**NEW QUESTION 476**

Which of the following will occur when an EC2 instance in a VPC with an associated Elastic IP is stopped and started? (Choose 2 answers)

- A. The Elastic IP will be dissociated from the instance
- B. All data on instance-store devices will be lost
- C. All data on EBS (Elastic Block Store) devices will be lost
- D. The ENI (Elastic Network Interface) is detached
- E. The underlying host for the instance is changed

**Answer:** BE

**NEW QUESTION 478**

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

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

**Answer:** A

**NEW QUESTION 480**

You are deploying an application on EC2 that must call AWS APIs. What method of securely passing credentials to the application should you use?

- A. Use AWS Identity and Access Management roles for EC2 instances.
- B. Pass API credentials to the instance using instance userdata.
- C. Embed the API credentials into your JAR files.
- D. Store API credentials as an object in Amazon Simple Storage Service.

**Answer:** A

**NEW QUESTION 481**

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 483**

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 486**

How can software determine the public and private IP addresses of the EC2 instance that it is running on?

- A. Query the local instance metadata.
- B. Query the local instance userdata.
- C. Query the appropriate Amazon CloudWatch metric.
- D. Use an ipconfig or ifconfig command.

**Answer:** A

**NEW QUESTION 489**

You have an application running in us-west-2 that requires six EC2 instances running at all times. With three AZs available in that region (us-west-2a, us-west-2b, and us-west-2c), which of the following deployments provides 100 percent fault tolerance if any single AZ in us-west-2 becomes unavailable? Choose 2 answers

- A. Us-west-2a with two EC2 instances, us-west-2b with two EC2 instances, and us-west-2c with two EC2 instances
- B. Us-west-2a with three EC2 instances, us-west-2b with three EC2 instances, and us-west-2c with no EC2 instances
- C. Us-west-2a with four EC2 instances, us-west-2b with two EC2 instances, and us-west-2c with two EC2 instances
- D. Us-west-2a with six EC2 instances, us-west-2b with six EC2 instances, and us-west-2c with no EC2 instances
- E. Us-west-2a with three EC2 instances, us-west-2b with three EC2 instances, and us-west-2c with three EC2 instances

**Answer:** DE

**Explanation:**

- option A : 2 2 2
- option B : 3 3 –
- option C : 4 2 2
- option D : 6 6 –
- option E : 3 3 3

so if one availability zone fails you need to have a backup of 6 instances running only D & E has that chance

**NEW QUESTION 490**

After creating a new AWS account, you use the API to request 40 on-demand EC2 instances in a single AZ. After 20 successful requests, subsequent requests failed. What could be a reason for this issue, and how would you resolve it?

- A. You encountered a soft limit of 20 instances per regio
- B. Submit the limit increase form and retry the failed requests once approved.
- C. AWS allows you to provision no more than 20 instances per Availability Zon
- D. Select a different Availability Zone and retry the failed request.
- E. You need to use Amazon Virtual Private Cloud (VPC) in order to provision more than 20 instances in a single Availability Zon
- F. Simply terminate the resources already provisioned and re-launch them all in a VPC.
- G. You encountered an API throttling situation and should try the failed requests using an exponential decay retry algorithm.

**Answer:** A

**NEW QUESTION 494**

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

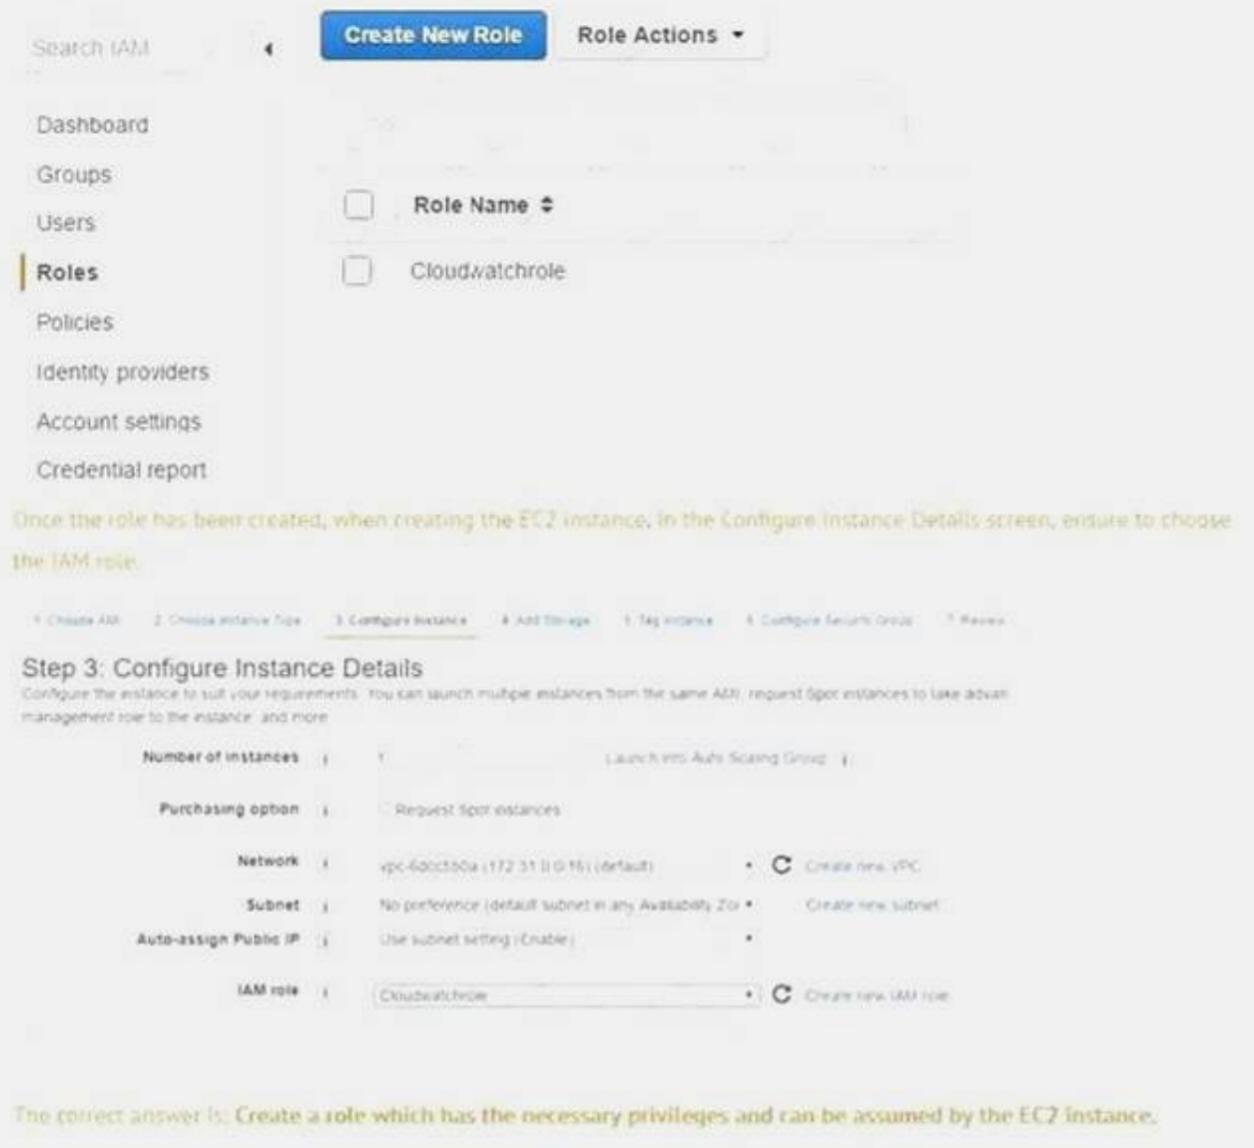
- A. Create a role which has the necessary and can be assumed by the EC2 instance.
- B. Use the API credentials from an EC2 instance.
- C. Use the API credentials from a bastion host.
- D. Use the API credentials from a NAT Instanc

**Answer:** A

**Explanation:**

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

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



**NEW QUESTION 497**

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

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

**Answer: B**

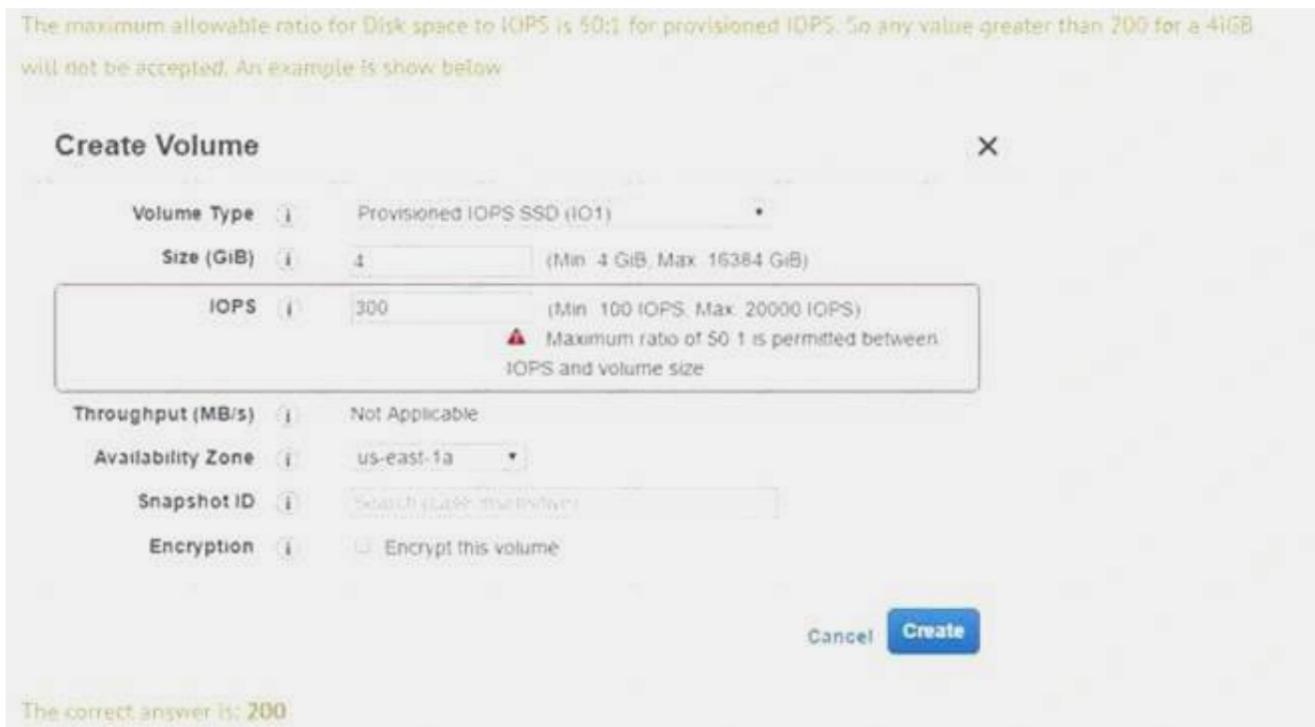
**NEW QUESTION 500**

If a provisioned IOPS volume of 4iGB is created, what are the possible correct values for IOPS for the volume in order for it to be created?

- A. 200
- B. 300
- C. 400
- D. 500

**Answer: A**

**Explanation:**



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

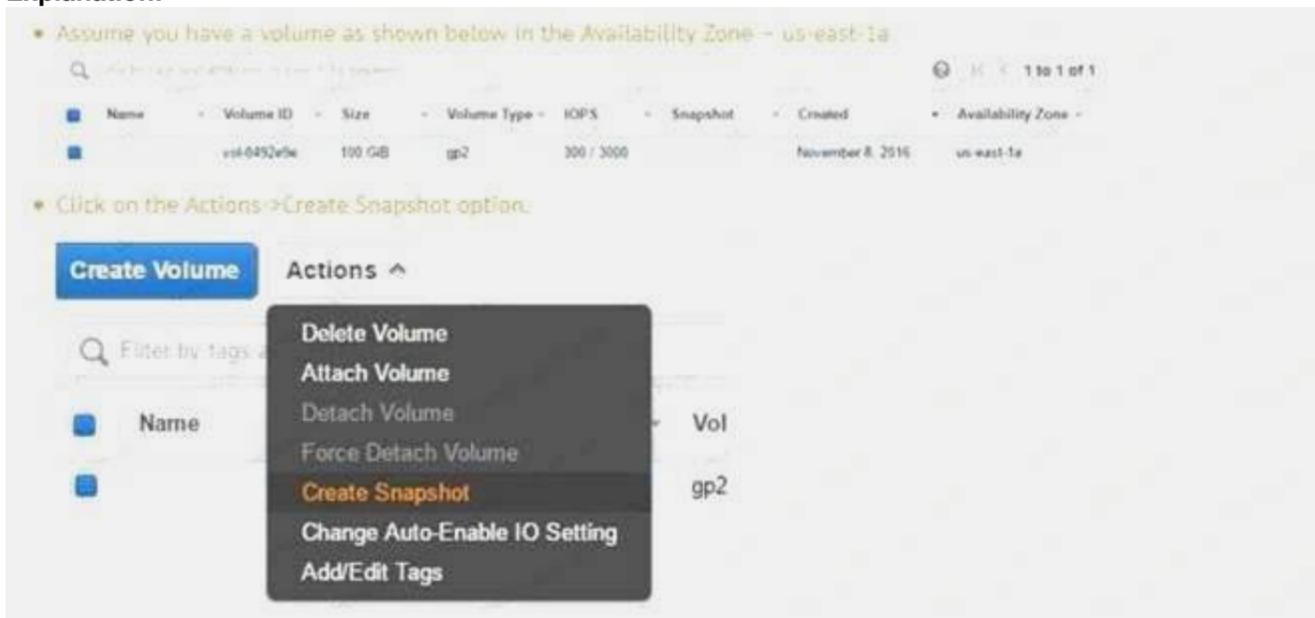
**NEW QUESTION 502**

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

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

**Answer: C**

**Explanation:**



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

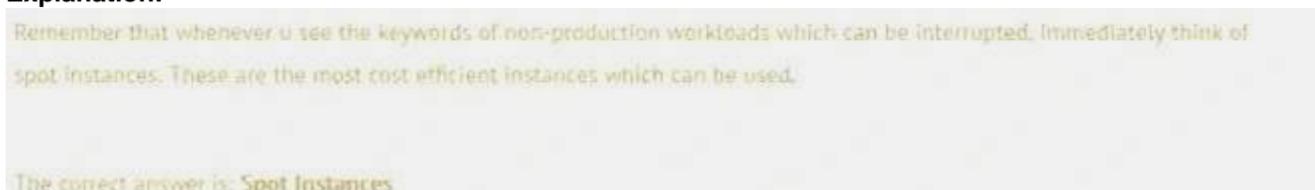
**NEW QUESTION 507**

A company is hosting EC2 instances which focuses on work-loads are on non-production and nonpriority batch loads. Also these processes can be interrupted at any time.  
What is the best pricing model which can be used for EC2 instances in ülis case?

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

**Answer: C**

**Explanation:**



**NEW QUESTION 510**

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

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

**Answer: A**

**NEW QUESTION 511**

You have been asked to design a fault-tolerant and scalable web application across three Availability Zones. The presentation logic will reside on web servers behind an ELB Classic Load Balancer, and the application logic will reside on a set of app servers behind a second load balancer. How should you use Auto Scaling groups?

- A. Deploy one Auto Scaling group that includes all the web and app servers across all Availability Zones.
- B. Deploy three Auto Scaling groups: one for each Availability Zone that includes both web and app servers.
- C. Deploy two Auto Scaling groups: one for the web servers in all Availability Zones and one for the app servers in all Availability Zones.
- D. Deploy six Auto Scaling groups: a web server group in each Availability Zone and an app server group in each Availability Zone.

**Answer: C**

**NEW QUESTION 514**

A company has a workflow that uploads video files from their data center to AWS for transcoding. They use Amazon EC2 worker instances that pull transcoding jobs from SQS.

Why is SQS an appropriate service for this scenario?

- A. SQS can accommodate message payloads of any size.
- B. SQS checks the health of the worker instances.
- C. SQS synchronously provides transcoding output.
- D. SQS decouples the transcoding task from the upload.

**Answer: D**

**NEW QUESTION 516**

You originally built a VPC for a two-tier application. The subnets for the web and data tiers use all the IP address space in the VPC. Now you want to add subnets for an application tier.

How can you accommodate the new subnets in your VPC?

- A. Reduce the CIDR block ranges of the existing subnets to make room for the new subnets.
- B. Build a new VPC that can accommodate all the subnets, and migrate the application to the new VPC.
- C. Change the CIDR block for the VPC to create enough free address space for the new subnets.
- D. Create the new subnets in the VPC; the VPC will automatically scale to accommodate the new subnets.

**Answer: A**

**NEW QUESTION 521**

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 523**

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

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

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

**Answer: A**

**NEW QUESTION 525**

A colleague asked for your advice about how to easily deploy, monitor, and scale a three-tier LAMP (Linux, Apache, MySQL, PHP) application on AWS. Your colleague has time and staffing constraints and wants to deploy and manage the application with minimal effort. Which AWS service would you suggest?

- A. Elastic Beanstalk
- B. Data Pipeline
- C. CloudFormation
- D. CodeDeploy

**Answer:** A

**NEW QUESTION 526**

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 529**

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 531**

.....

## 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>