

Exam Questions SAA-C01

AWS Certified Solutions Architect - Associate

<https://www.2passeasy.com/dumps/SAA-C01/>



NEW QUESTION 1

A 3-tier e-commerce web application is current deployed on-premises and will be migrated to AWS for greater scalability and elasticity. The web server currently shares read-only data using a network distributed file system The app server tier uses a clustering mechanism for discovery and shared session state that depends on IP multicast The database tier uses shared-storage clustering to provide database fail over capability, and uses several read slaves for scaling Data on all servers and the distributed file system directory is backed up weekly to off-site tapes.

Which AWS storage and database architecture meets the requirements of the application?

- A. Web servers: store read-only data in S3, and copy from S3 to root volume at boot tim
- B. App servers: share state using a combination of DynamoDB and IP unicas
- C. Database: use RDS with multi- AZ deployment and one or more read replica
- D. Backup: web servers, app servers, and database backed up weekly to Glacier using snapshots.
- E. Web servers: store read-only data in an EC2 NFS server, mount to each web server at boot tim
- F. App servers: share state using a combination of DynamoDB and IP multicas
- G. Database: use RDS with multi-AZ deployment and one or more Read Replica
- H. Backup: web and app servers backed up weekly via AMIs, database backed up via DB snapshots.
- I. Web servers: store read-only data in S3, and copy from S3 to root volume at boot tim
- J. App servers: share state using a combination of DynamoDB and IP unicas
- K. Database: use RDS with multi- AZ deployment and one or more Read Replica
- L. Backup: web and app servers backed up weekly via AMIs, database backed up via DB snapshots.
- M. Web servers: store read-only data in S3, and copy from S3 to root volume at boot tim
- N. App servers: share state using a combination of DynamoDB and IP unicas
- O. Database: use RDS with multi- AZ deploymen
- P. Backup: web and app servers backed up weekly via AMIs, database backed up via DB snapshots.

Answer: A

Explanation:

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

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

NEW QUESTION 2

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

- A. Route 53 Record Sets
- B. IM1 Roles
- C. Elastic IP Addresses (EIP)
- D. EC2 Key Pairs
- E. Launch configurations
- F. Security Groups

Answer: AB

Explanation:

The Route 53 and IAM are global.

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

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

NEW QUESTION 3

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

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

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

- A. Yes, you should deploy two Memcached ElastiCache Clusters in different AZs because the RDS instance will not be able to handle the load if the cache node fails.
- B. No, if the cache node fails you can always get the same data from the DB withouthaving any availability impact.
- C. No, if the cache node fails the automated ElastiCache node recovery feature will prevent any availability impact.
- D. Yes, you should deploy the Memcached ElastiCache Cluster with two nodes in the same AZ as the RDS DB master instance to handle the load if one cache node fails.

Answer: A

Explanation:

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

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

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

To mitigate the impact of an availability zone failure, locate your nodes in as many availability zones as possible. In the unlikely event of an AZ failure, you will lose

only the data cached in that AZ, not the data cached in the other AZs.

NEW QUESTION 4

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

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

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

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

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

Answer: C

Explanation:

http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_ecommerce_checkout_13.pdf

NEW QUESTION 5

Your company hosts a social media site supporting users in multiple countries. You have been asked to provide a highly available design for the application that leverages multiple regions for the most recently accessed content and latency sensitive portions of the website. The most latency sensitive component of the application involves reading user preferences to support web site personalization and ad selection.

In addition to running your application in multiple regions, which option will support this application's requirements?

- A. Serve user content from S3. CloudFront and use Route53 latency-based routing between ELBs in each region. Retrieve user preferences from a local DynamoDB table in each region and leverage SQS to capture changes to user preferences with SNS workers for propagating updates to each table.
- B. Use the S3 Copy API to copy recently accessed content to multiple regions and serve user content from S3. CloudFront with dynamic content and an ELB in each region. Retrieve user preferences from an ElasticCache cluster in each region and leverage SNS notifications to propagate user preference changes to a worker node in each region.
- C. Use the S3 Copy API to copy recently accessed content to multiple regions and serve user content from S3. CloudFront and Route53 latency-based routing between ELBs. In each region, retrieve user preferences from a DynamoDB table and leverage SQS to capture changes to user preferences with SNS workers for propagating DynamoDB updates.
- D. Serve user content from S3. CloudFront with dynamic content, and an ELB in each region. Retrieve user preferences from an ElasticCache cluster in each region and leverage Simple Workflow (SWF) to manage the propagation of user preferences from a centralized DB to each ElasticCache cluster.

Answer: A

Explanation:

http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_mediassharing_09.pdf

http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_adserving_06.pdf

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

A web company is looking to implement an external payment service into their highly available application deployed in a VPC. Their application EC2 instances are behind a public-facing ELB. Auto scaling is used to add additional instances as traffic increases. Under normal load the application runs 2 instances in the Auto Scaling group but at peak it can scale 3x in size. The application instances need to communicate with the payment service over the Internet which requires whitelisting of all public IP addresses used to communicate with it. A maximum of 4 whitelisting IP addresses is allowed at a time and can be added through an API.

How should they architect their solution?

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

Answer: A

Explanation:

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

NEW QUESTION 8

You have deployed a three-tier web application in a VPC with a CIDR block of 10.0.0.0/28. You initially deploy two web servers, two application servers, two database servers and one NAT instance for a total of seven EC2 instances. The web, application and database servers are deployed across two availability zones (AZs). You also deploy an ELB in front of the two web servers, and use Route53 for DNS. Web traffic gradually increases in the first few days following the deployment, so you attempt to double the number of instances in each tier of the application to handle the new load. Unfortunately, some of these new instances fail to launch. Which of the following could be the root cause? (Choose two.)

- A. AWS reserves the first and the last private IP address in each subnet's CIDR block so you do not have enough addresses left to launch all of the new EC2 instances
- B. The Internet Gateway (IGW) of your VPC has scaled up, adding more instances to handle the traffic spike, reducing the number of available private IP addresses for new instance launches
- C. The ELB has scaled up, adding more instances to handle the traffic spike, reducing the number of available private IP addresses for new instance launches
- D. AWS reserves one IP address in each subnet's CIDR block for Route53 so you do not have enough addresses left to launch all of the new EC2 instances
- E. AWS reserves the first four and the last IP address in each subnet's CIDR block so you do not have enough addresses left to launch all of the new EC2 instances

Answer: CE

Explanation:

http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

The first four IP addresses and the last IP address in each subnet CIDR block are not available for you to use, and cannot be assigned to an instance. For example, in a subnet with CIDR block 10.0.0.0/24, the following five IP addresses are reserved:

- 10.0.0.0: Network address.
- 10.0.0.1: Reserved by AWS for the VPC router.
- 10.0.0.2: Reserved by AWS. The IP address of the DNS server is always the base of the VPC network range plus two; however, we also reserve the base of each subnet range plus two. For more information, see Amazon DNS Server.
- 10.0.0.3: Reserved by AWS for future use.
- 10.0.0.255: Network broadcast address. We do not support broadcast in a VPC, therefore we reserve this address.

NEW QUESTION 9

You require the ability to analyze a large amount of data, which is stored on Amazon S3 using Amazon Elastic Map Reduce. You are using the cc2.8xlarge instance type, whose CPUs are mostly idle during processing. Which of the below would be the most cost-efficient way to reduce the runtime of the job?

- A. Create smaller files on Amazon S3.
- B. Add additional cc2.8xlarge instances by introducing a task group.
- C. Use smaller instances that have higher aggregate I/O performance.
- D. Create fewer, larger files on Amazon S3.

Answer: C

NEW QUESTION 10

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 and 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 days
- 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 days
- 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 days
- 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 days
- L. CloudFront to serve HLS transcoded videos from Glacier.

Answer: C

NEW QUESTION 10

You are designing a data leak prevention solution for your VPC environment. You want your VPC Instances to be able to access software depots and distributions on the Internet for product updates. The depots and distributions are accessible via third party CONs by their URLs. You want to explicitly deny any other outbound connections from your VPC instances to hosts on the Internet.

Which of the following options would you consider?

- A. Configure a web proxy server in your VPC and enforce URL-based rules for outbound access Remove default routes.
- B. Implement security groups and configure outbound rules to only permit traffic to software depots.
- C. Move all your instances into private VPC subnets remove default routes from all routing tables and add specific routes to the software depots and distributions only.
- D. Implement network access control lists to all specific destinations, with an Implicit deny as a rule

Answer: A

Explanation:

Organizations usually implement proxy solutions to provide URL and web content filtering, IDS/IPS, data loss prevention, monitoring, and advanced threat protection.

https://d0.awsstatic.com/aws-answers/Controlling_VPC_Egress_Traffic.pdf

NEW QUESTION 13

An administrator is using Amazon CloudFormation to deploy a three tier web application that consists of a web tier and application tier that will utilize Amazon DynamoDB for storage when creating the CloudFormation template which of the following would allow the application instance access to the DynamoDB tables without exposing API credentials?

- A. Create an Identity and Access Management Role that has the required permissions to read and write from the required DynamoDB table and associate the Role to the application instances by referencing an instance profile.
- B. Use the Parameter section in the Cloud Formation template to have the user input Access and Secret Keys from an already created IAM user that has the permissions required to read and write from the required DynamoDB table.
- C. Create an Identity and Access Management Role that has the required permissions to read and write from the required DynamoDB table and reference the Role in the instance profile property of the application instance.
- D. Create an identity and Access Management user in the CloudFormation template that has permissions to read and write from the required DynamoDB table, use the GetAtt function to retrieve the Access and secret keys and pass them to the application instance through user-data.

Answer: C

NEW QUESTION 16

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

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

Answer: C

NEW QUESTION 19

You have an application running on an EC2 Instance which will allow users to download files from a private S3 bucket using a pre-assigned URL. Before generating the URL the application should verify the existence of the file in S3.

How should the application use AWS credentials to access the S3 bucket securely?

- A. Use the AWS account access Keys the application retrieves the credentials from the source code of the application.
- B. Create an IAM user for the application with permissions that allow list access to the S3 bucket launch the instance as the IAM user and retrieve the IAM user's credentials from the EC2 instance user data.
- C. Create an IAM role for EC2 that allows list access to objects in the S3 bucket
- D. Launch the instance with the role, and retrieve the role's credentials from the EC2 Instance metadata
- E. Create an IAM user for the application with permissions that allow list access to the S3 bucket
- F. The application retrieves the IAM user credentials from a temporary directory with permissions that allow read access only to the application user

Answer: C

Explanation:

Reference

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-metadata.html>

NEW QUESTION 22

Your company policies require encryption of sensitive data at rest. You are considering the possible options for protecting data while storing it at rest on an EBS data volume, attached to an EC2 instance. Which of these options would allow you to encrypt your data at rest? (Choose three.)

- A. Implement third party volume encryption tools
- B. Do nothing as EBS volumes are encrypted by default
- C. Encrypt data inside your applications before storing it on EBS
- D. Encrypt data using native data encryption drivers at the file system level
- E. Implement SSL/TLS for all services running on the server

Answer: ACD

Explanation:

Not E since SSL/TLS is encryption in transfer (https) and not encryption of sensitive data at rest. And B is just not true. Although you nowadays can add encryption when creating a EBS volume but it is NOT turned on by default.

NEW QUESTION 25

You have a periodic Image analysis application that gets some files. In Input analyzes them and for each file writes some data in output to a ten file the number of files in input per day is high and concentrated in a few hours of the day.

Currently you have a server on EC2 with a large EBS volume that hosts the input data and the results it takes almost 20 hours per day to complete the process What services could be used to reduce the elaboration time and improve the availability of the solution?

- A. S3 to store I/O file
- B. SQS to distribute elaboration commands to a group of hosts working in paralle
- C. Auto scaling to dynamically size the group of hosts depending on the length of the SQS queue
- D. EBS with Provisioned IOPS (PIOPS) to store I/O file
- E. SNS to distribute elaboration commands to a group of hosts working in parallel Auto Scaling to dynamically size the group of hosts depending on the number of SNS notifications
- F. S3 to store I/O files, SNS to distribute evaporation commands to a group of hosts working in paralle
- G. Auto scaling to dynamically size the group of hosts depending on the number of SNS notifications
- H. EBS with Provisioned IOPS (PIOPS) to store I/O files SOS to distribute elaboration commands to a group of hosts working in parallel Auto Scaling to dynamically size the group ot hosts depending on the length of the SQS queue.

Answer: D

Explanation:

Amazon EBS allows you to create storage volumes and attach them to Amazon EC2 instances. Once attached, you can create a file system on top of these volumes, run a database, or use them in any other way you would use a block device. Amazon EBS volumes are placed in a specific Availability Zone, where they are automatically replicated to protect you from the failure of a single component. Amazon EBS provides three volume types: General Purpose (SSD), Provisioned IOPS (SSD), and Magnetic. The three volume types differ in performance characteristics and cost, so you can choose the right storage performance and price for the needs of your applications. All EBS volume types offer the same durable snapshot capabilities and are designed for 99.999% availability.

NEW QUESTION 30

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

- A. Log clicks in weblogs by URL store to Amazon S3, and then analyze with Elastic MapReduce
- B. Push web clicks by session to Amazon Kinesis and analyze behavior using Kinesis workers
- C. Write click events directly to Amazon Redshift and then analyze with SQL
- D. Publish web clicks by session to an Amazon SQS queue men periodically drain these events to Amazon RDS and analyze with sol

Answer: B

NEW QUESTION 35

An AWS customer runs a public blogging website. The site users upload two million blog entries a month. The average blog entry size is 200 KB. The access rate to blog entries drops to negligible 6 months after publication and users rarely access a blog entry 1 year after publication. Additionally, blog entries have a high update rate during the first 3 months following publication, this drops to no updates after 6 months. The customer wants to use CloudFront to improve his user's load times. Which of the following recommendations would you make to the customer?

- A. Duplicate entries into two different buckets and create two separate CloudFront distributions where S3 access is restricted only to Cloud Front identity
- B. Create a CloudFront distribution with "US'Europe price class for US/Europe users and a different CloudFront distribution with All Edge Locations' for the remaining users.
- C. Create a CloudFront distribution with S3 access restricted only to the CloudFront identity and partition the blog entry's location in S3 according to the month it was uploaded to be used with CloudFront behaviors.
- D. Create a CloudFronl distribution with Restrict Viewer Access Forward Query string set to true and minimum TTL of 0.

Answer: C

NEW QUESTION 36

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 38

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 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 tiles 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 41

You are running a successful multitier web application on AWS and your marketing department has asked you to add a reporting tier to the application. The reporting tier will aggregate and publish status reports every 30 minutes from user-generated information that is being stored in your web application's database. You are currently running a Multi-AZ RDS MySQL instance for the database tier. You also have implemented ElastiCache as a database caching layer between the application tier and database tier. Please select the answer that will allow you to successfully implement the reporting tier with as little impact as possible to your database.

- A. Continually send transaction logs from your master database to an S3 bucket and generate the reports off the S3 bucket using S3 byte range requests.
- B. Generate the reports by querying the synchronously replicated standby RDS MySQL instance maintained through Multi-AZ.
- C. Launch a RDS Read Replica connected to your Multi AZ master database and generate reports by querying the Read Replica.
- D. Generate the reports by querying the ElastiCache database caching tier

Answer: C

Explanation:

Amazon RDS allows you to use read replicas with Multi-AZ deployments. In Multi-AZ deployments for MySQL, Oracle, SQL Server, and PostgreSQL, the data in your primary DB Instance is synchronously replicated to a standby instance in a different Availability Zone (AZ). Because of their synchronous replication, Multi-AZ deployments for these engines offer greater data durability benefits than do read replicas. (In all Amazon RDS for Aurora deployments, your data is automatically replicated across 3 Availability Zones.)

You can use Multi-AZ deployments and read replicas in conjunction to enjoy the complementary benefits of each. You can simply specify that a given Multi-AZ deployment is the source DB Instance for your Read replicas. That way you gain both the data durability and availability benefits of Multi-AZ deployments and the read scaling benefits of read replicas.

Note that for Multi-AZ deployments, you have the option to create your read replica in an AZ other than that of the primary and the standby for even more redundancy. You can identify the AZ corresponding to your standby by looking at the "Secondary Zone" field of your DB Instance in the AWS Management Console.

NEW QUESTION 45

You must assign each server to at least _____ security group

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

Answer: D

Explanation:

Your AWS account automatically has a default security group per region for EC2-Classic. When you create a VPC, we automatically create a default security group for the VPC. If you don't specify a different security group when you launch an instance, the instance is automatically associated with the appropriate default security group. <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html>

NEW QUESTION 46

What does RRS stand for when talking about S3?

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

Answer: D

Explanation:

In Amazon S3, RRS stands for Reduced Redundancy Storage. Reduced redundancy storage stores objects on multiple devices across multiple facilities, providing 400 times the durability of a typical disk drive, but it does not replicate objects as many times as Amazon S3 standard storage. In addition, reduced redundancy storage is designed to sustain the loss of data in a single facility. <http://docs.aws.amazon.com/AmazonS3/latest/dev/UsingRRS.html>

NEW QUESTION 47

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

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

Answer: C

NEW QUESTION 48

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 51

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

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

Answer: C

NEW QUESTION 53

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

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

Answer: C

NEW QUESTION 54

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

- A. TRUE
- B. FALSE

Answer: A

Explanation:

AWS Direct Connect Keys

AWS Direct Connect implements the following policy keys:

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

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

Note

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

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

http://docs.aws.amazon.com/directconnect/latest/UserGuide/using_iam.html

NEW QUESTION 55

Amazon RDS automated backups and DB Snapshots are currently supported for only the _ storage engine

- A. InnoDB
- B. MyISAM

Answer: A

NEW QUESTION 59

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

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

Answer: C

Explanation:

<http://docs.aws.amazon.com/cli/latest/reference/ec2/create-snapshot.html>

NEW QUESTION 62

Typically, you want to check your application whether a request generated an error before you spend any time processing results. The easiest way to find out if an error occurred is to look for an ____ node in the response from the Amazon RDS API.

- A. Incorrect
- B. Error
- C. FALSE

Answer: B

Explanation:

Typically, you want your application to check whether a request generated an error before you spend any time processing results. The easiest way to find out if an error occurred is to look for an Error node in the response from the Amazon RDS API.

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

NEW QUESTION 64

What are the two permission types used by AWS?

- A. Resource-based and Product-based
- B. Product-based and Service-based
- C. Service-based
- D. User-based and Resource-based

Answer: D

NEW QUESTION 69

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

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

Answer: B

NEW QUESTION 74

You must increase storage size in increments of at least _____ %

- A. 40
- B. 20
- C. 50
- D. 10

Answer: D

NEW QUESTION 79

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 81

What does a "Domain" refer to in Amazon SWF?

- A. A security group in which only tasks inside can communicate with each other
- B. A special type of worker
- C. A collection of related Workflows
- D. The DNS record for the Amazon SWF service

Answer: C

Explanation:

Domains provide a way of scoping Amazon SWF resources within your AWS account. All the components of a workflow, such as the workflow type and activity types, must be specified to be in a domain. It is possible to have more than one workflow in a domain; however, workflows in different domains cannot interact with each other. <http://docs.aws.amazon.com/amazonswf/latest/developerguide/swf-dev-domain.html>

NEW QUESTION 84

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

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

Answer: B

Explanation:

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

NEW QUESTION 87

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_fc5fbcd3-0405-3b7d-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 88

What does the following command do with respect to the Amazon EC2 security groups? `ec2-revoke RevokeSecurityGroupIngress`

- A. Removes one or more security groups from a rule.
- B. Removes one or more security groups from an Amazon EC2 instance.
- C. Removes one or more rules from a security group.
- D. Removes a security group from our account

Answer: C

Explanation:

Removes one or more ingress rules from a security group. The values that you specify in the revoke request (for example, ports) must match the existing rule's values for the rule to be removed. <http://docs.aws.amazon.com/cli/latest/reference/ec2/revoke-security-group-ingress.html>

revoke-security-group-ingress

Note:

To specify multiple rules in a single command use the `--ip-permissions` option.

Description

Removes one or more ingress rules from a security group. The values that you specify in the revoke request (for example, ports) must match the existing rule's values for the rule to be removed.

Each rule consists of the protocol and the CIDR range or source security group. For the TCP and UDP protocols, you must also specify the destination port or range of ports. For the ICMP protocol, you must also specify the ICMP type and code.

Rule changes are propagated to instances within the security group as quickly as possible. However, a small delay might occur.

NEW QUESTION 92

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

- A. No
- B. Yes

Answer: A

NEW QUESTION 94

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

- A. TRUE
- B. FALSE

Answer: B

Explanation:

If you choose not to create a final DB snapshot, you will not be able to later restore the DB instance to its final state. When you delete a DB instance, all automated backups are deleted and cannot be recovered. Manual DB snapshots of the instance are not deleted.

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_DeleteInstance.html

NEW QUESTION 99

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

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

Answer: B

NEW QUESTION 100

How many types of block devices does Amazon EC2 support?

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

Answer: A

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/block-device-mapping-concepts.html> Amazon EC2 supports two types of block devices:

Instance store volumes (virtual devices whose underlying hardware is physically attached to the host computer for the instance)

EBS volumes (remote storage devices)

A block device mapping defines the block devices (instance store volumes and EBS volumes) to attach to an instance.

Block Device Mapping Concepts

A *block device* is a storage device that moves data in sequences of bytes or bits (blocks). These devices support random access and generally use buffered I/O. Examples include hard disks, CD-ROM drives, and flash drives. A block device can be physically attached to a computer or accessed remotely as if it were physically attached to the computer. Amazon EC2 supports **two types** of block devices:

- Instance store volumes (virtual devices whose underlying hardware is physically attached to the host computer for the instance)
- EBS volumes (remote storage devices)

NEW QUESTION 105

Provisioned IOPS Costs: you are charged for the IOPS and storage whether or not you use them in a given month.

- A. FALSE
- B. TRUE

Answer: B

Explanation:

Volume storage for EBS Provisioned IOPS SSD (io1) volumes is charged by the amount you provision in GB per month, until you release the storage. With Provisioned IOPS SSD (io1) volumes, you are also charged by the amount you provision in IOPS (input/output operations per second) multiplied by the percentage of days you provision for the month. For example, if you provision a volume with 1000 IOPS, and keep this volume for 15 days in a 30 day month, then in a Region that charges \$0.10 per provisioned IOPS-month, you would be charged \$50 for the IOPS that you provision (\$0.10 per provisioned IOPS-month * 1000 IOPS provisioned * 15 days/30). You will be charged for the IOPS provisioned on a volume even when the volume is detached from an instance.

NEW QUESTION 107

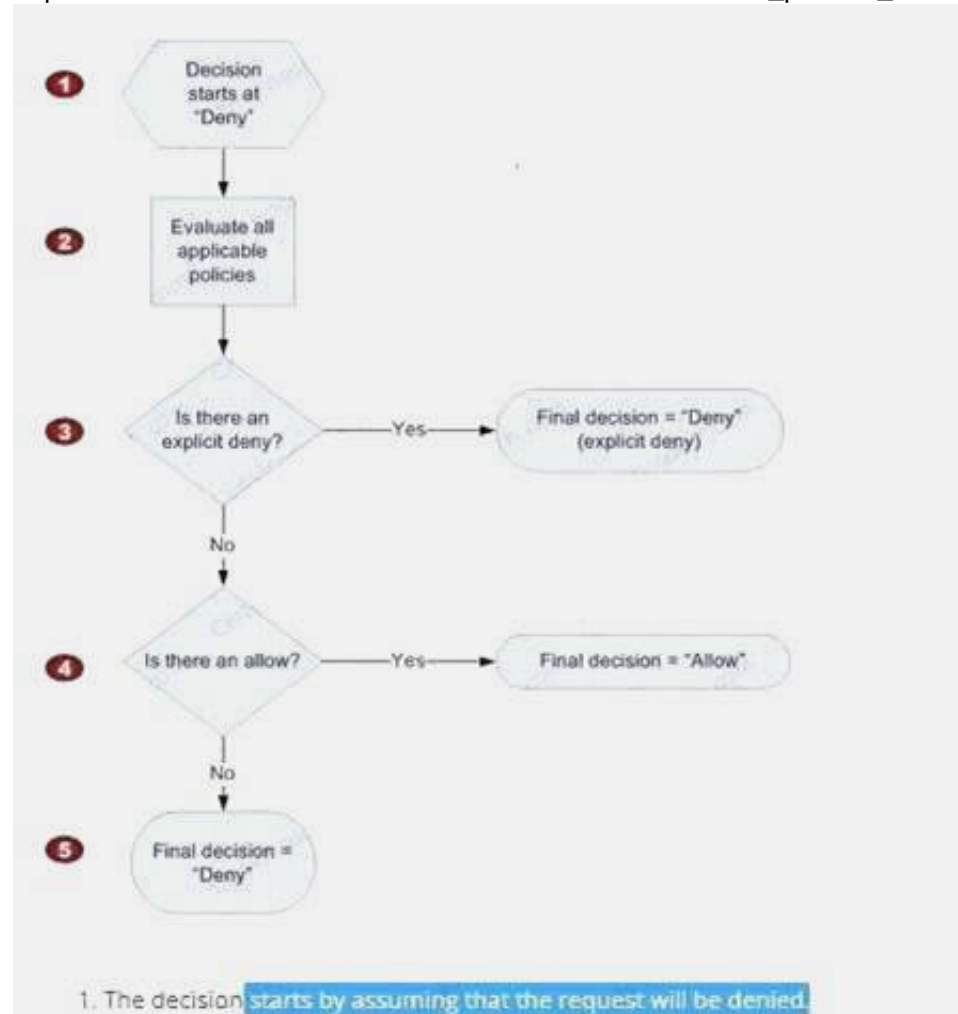
IAM's Policy Evaluation Logic always starts with a default ____ for every request, except for those that use the AWS account's root security credentials b

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

Answer: B

Explanation:

http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_evaluation-logic.html



NEW QUESTION 109

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

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

Answer: D

Explanation:

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

Viewing Volume Information

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

To view information about an EBS volume using the console

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

NEW QUESTION 114

While creating the snapshots using the API, which Action should I be using?

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

Answer: D

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/CommandLineReference/ApiReference-cmd-CreateSnapshot.html>

NEW QUESTION 117

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 118

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

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

Answer: B

NEW QUESTION 120

Are Reserved Instances available for Multi-AZ Deployments?

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

Answer: B

Explanation:

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

NEW QUESTION 121

Amazon RDS automated backups and DB Snapshots are currently supported for only the _____ storage engine

- A. MyISAM
- B. InnoDB

Answer: B

NEW QUESTION 124

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

NEW QUESTION 128

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 130

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 133

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

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

Answer: C

NEW QUESTION 136

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

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

Answer: D

Explanation:

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

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_Tagging.html <http://docs.aws.amazon.com/cli/latest/reference/rds/list-tags-for-resource.html>

NEW QUESTION 140

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 143

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 144

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

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

Answer: B

Explanation:

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

NEW QUESTION 146

When automatic failover occurs, Amazon RDS will emit a DB Instance event to inform you that automatic failover occurred. You can use the ____ to return information about events related to your DB Instance

- A. FetchFailure
- B. DescribeFailure
- C. DescribeEvents
- D. FetchEvents

Answer: C

Explanation:

Q: Will I be alerted when automatic failover occurs?

Yes, Amazon RDS will emit a DB Instance event to inform you that automatic failover occurred. You can use the DescribeEvents to return information about events related to your DB Instance, or click the "DB Events" section of the AWS Management Console

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

NEW QUESTION 150

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 154

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

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

Answer: A

Explanation:

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

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

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

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

Default Security Groups

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

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

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

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

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

NEW QUESTION 157

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

NEW QUESTION 162

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

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

Answer: B

Explanation:

"If an EBS volume is the root device of an instance, you must stop the instance before you can detach the volume."
<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-detaching-volume.html>

NEW QUESTION 163

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 164

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

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

Answer: C

Explanation:

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

NEW QUESTION 166

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 167

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

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

Answer: B

NEW QUESTION 171

Is it possible to access your EBS snapshots?

- A. Yes, through the Amazon S3 APIs.
- B. Yes, through the Amazon EC2 APIs.
- C. No, EBS snapshots cannot be accessed; they can only be used to create a new EBS volume.
- D. EBS doesn't provide snapshot

Answer: B

Explanation:

https://aws.amazon.com/ebs/faqs/?nc1=h_ls

Q: Will I be able to access my snapshots using the regular Amazon S3 API? No, snapshots are only available through the Amazon EC2 API.

NEW QUESTION 173

Does Route 53 support MX Records?

- A. Yes.
- B. It supports CNAME records, but not MX records.
- C. No
- D. Only Primary MX record
- E. Secondary MX records are not supported.

Answer: A

Explanation:

<http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/ResourceRecordTypes.html#MXForm>

MX Format

Each value for an MX resource record set actually contains two values:

- An integer that represents the priority for an email server
- The domain name of the email server

If you specify only one server, the priority can be any integer between 0 and 65535. If you specify multiple servers, the value that you specify for the priority indicates which email server you want email to be routed to first, second, and so on. For example, if you have two email servers and you specify values of 10 and 20 for the priority, email always goes to the server with a priority of 10 unless it's unavailable. If you specify values of 10 and 10, email is routed to the two servers approximately equally.

Example for the Amazon Route 53 console

```
10 mail.example.com
```

Example for the Amazon Route 53 API

```
<Value>10 mail.example.com</Value>
```

NEW QUESTION 177

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 182

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

- A. Select public for all the accounts and check mark those accounts with whom you want to expose the snapshots and click save.
- B. SelectPrivate, enter the IDs of those AWS accounts, and clickSave.
- C. SelectPublic, enter the IDs of those AWS accounts, and clickSave.
- D. SelectPublic, mark the IDs of those AWS accounts as private, and clickSav

Answer: B

Explanation:

"To expose the snapshot to only specific AWS accounts, choose Private, enter the ID of the AWS account (without hyphens) in the AWS Account Number field, and choose Add Permission. Repeat until you've added all the required AWS accounts" <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-modifying-snapshot-permissions.html>

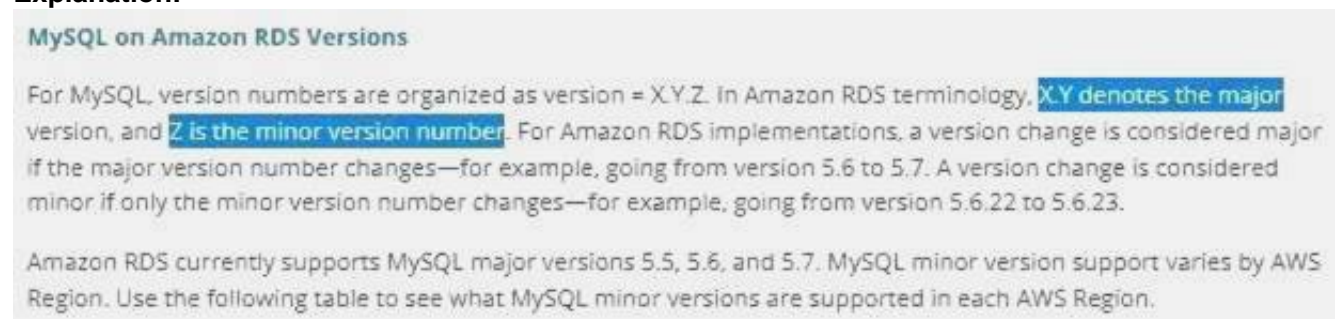
NEW QUESTION 184

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 187

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

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

Answer: A

NEW QUESTION 192

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

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

Answer: B

NEW QUESTION 197

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

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

C. Yes

Answer: C

NEW QUESTION 201

What does Amazon RDS stand for?

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

Answer: B

NEW QUESTION 202

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

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

Answer: A

NEW QUESTION 207

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

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

Answer: D

NEW QUESTION 208

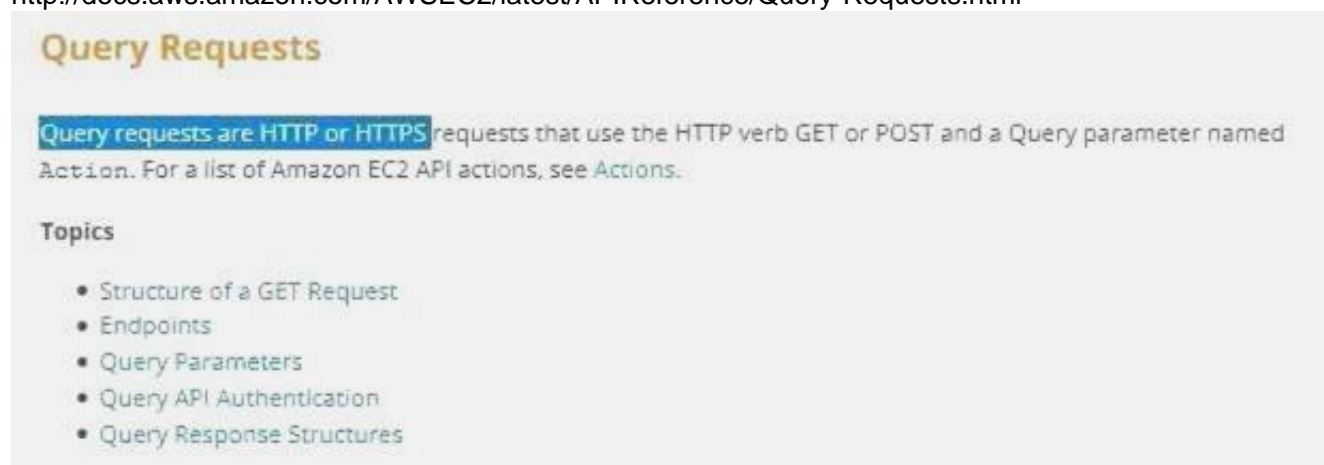
REST or Query requests are HTTP or HTTPS requests that use an HTTP verb (such as GET or POST) and a parameter named Action or Operation that specifies the API you are calling.

- A. FALSE
- B. TRUE

Answer: B

Explanation:

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



NEW QUESTION 213

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

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

Answer: C

Explanation:

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

NEW QUESTION 216

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.DBInstanceClass.Previous>

NEW QUESTION 220

Will I be charged if the DB instance is idle?

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

Answer: B

NEW QUESTION 225

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

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

Answer: C

NEW QUESTION 226

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 228

Amazon RDS supports SOAP only through _____ .

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

Answer: D

Explanation:

Amazon RDS supports SOAP only through HTTPS

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

WSDL and Schema Definitions

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

Note

Amazon RDS supports SOAP only through HTTPS.

NEW QUESTION 233

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 236


HTTP Query-based requests are HTTP requests that use the HTTP verb GET or POST and a Query parameter named ____ .

- A. Action
- B. Value
- C. Reset
- D. Retrieve

Answer: A

Explanation:

<http://docs.aws.amazon.com/AmazonS3/latest/dev/using-with-s3-actions.html>



NEW QUESTION 238

_____ embodies the "share-nothing" architecture and essentially involves breaking a large database into several smaller databases. Common ways to split a database include 1) splitting tables that are not joined in the same query onto different hosts or 2) duplicating a table across multiple hosts and then using a hashing algorithm to determine which host receives a given update.

- A. Sharding
- B. Failure recovery
- C. Federation
- D. DDL operations

Answer: A

Explanation:

Sharding embodies the "share-nothing" architecture and essentially just involves breaking a larger database up into smaller databases. Common ways to split a database are:

Splitting tables that are not joined in the same query onto different hosts Duplicating a table across multiple hosts and then splitting where a row goes.

More detailed information on the pros and cons of sharing can be found at the following sites: <http://technoroy.blogspot.com/2008/07/shard-database-design.html>

<http://www.hibernate.org/subprojects/shards.html>

How Amazon RDS Helps With Sharing Maintenance Overhead

NEW QUESTION 240

Does Amazon Route 53 support NS Records?

- A. Yes, it supports Name Service records.
- B. No
- C. It supports only MX records.
- D. Yes, it supports Name Server record

Answer: D

Explanation:

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

NEW QUESTION 242

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

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 246

When using consolidated billing there are two account types. What are they?

- A. Paying account and Linked account
- B. Parent account and Child account
- C. Main account and Sub account.
- D. Main account and Secondary account

Answer: A

Explanation:

You sign up for Consolidated Billing in the AWS Billing and Cost Management console, and designate your account as a payer account. Now your account can pay the charges of the other accounts, which are called linked accounts. The payer account and the accounts linked to it are called a Consolidated Billing account family. Source: <http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/consolidated-billing.html>

NEW QUESTION 247

A ____ is a document that provides a formal statement of one or more permissions.

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

Answer: A

Explanation:

http://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies.html

NEW QUESTION 252

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 253

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

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

Answer: C

Explanation:

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

NEW QUESTION 256

When using the following AWS services, which should be implemented in multiple Availability Zones for high availability solutions? Choose 2 answers

- A. Amazon DynamoDB
- B. Amazon Elastic Compute Cloud (EC2)
- C. Amazon Elastic Load Balancing
- D. Amazon Simple Notification Service (SNS)
- E. Amazon Simple Storage Service (S3)

Answer: BC

NEW QUESTION 261

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 265

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

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

Answer: BE

NEW QUESTION 270

Your application provides data transformation services. Files containing data to be transformed are first uploaded to Amazon S3 and then transformed by a fleet of spot EC2 instances. Files submitted by your premium customers must be transformed with the highest priority. How should you implement such a system?

- A. Use a DynamoDB table with an attribute defining the priority level
- B. Transformation instances will scan the table for tasks, sorting the results by priority level.
- C. Use Route 53 latency based-routing to send high priority tasks to the closest transformation instances.
- D. Use two SQS queues, one for high priority messages, the other for default priority
- E. Transformation instances first poll the high priority queue; if there is no message, they poll the default priority queue.
- F. Use a single SQS queue
- G. Each message contains the priority level
- H. Transformation instances poll high-priority messages first.

Answer: C

NEW QUESTION 273

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 IP

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 278

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

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

Answer: C

Explanation:

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

NEW QUESTION 279

Your web application front end consists of multiple EC2 instances behind an Elastic Load Balancer. You configured ELB to perform health checks on these EC2 instances, if an instance fails to pass health checks, which statement will be true?

- A. The instance gets terminated automatically by the ELB
- B. The instance gets quarantined by the ELB for root cause analysis.
- C. The instance is replaced automatically by the ELB
- D. The ELB stops sending traffic to the instance that failed its health check

Answer:

D

NEW QUESTION 280

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 285

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

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

Answer: BC

Explanation:

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

NEW QUESTION 288

A customer wants to leverage Amazon Simple Storage Service (S3) and Amazon Glacier as part of their backup and archive infrastructure. The customer plans to use third-party software to support this integration. Which approach will limit the access of the third party software to only the Amazon S3 bucket named "company-backup"?

- A. A custom bucket policy limited to the Amazon S3 API in the Amazon Glacier archive "companybackup"
- B. A custom bucket policy limited to the Amazon S3 API in "company-backup"
- C. A custom IAM user policy limited to the Amazon S3 API for the Amazon Glacier archive "companybackup".
- D. A custom IAM user policy limited to the Amazon S3 API in "company-backup".

Answer: D

Explanation:

<http://docs.aws.amazon.com/AmazonS3/latest/dev/example-policies-s3.html>

NEW QUESTION 293

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 <http://docs.aws.amazon.com/amazonglacier/latest/dev/uploading-archive-mpu.html>

NEW QUESTION 294

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

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

Answer: B

Explanation:

Elastic Load Balancing access logs

The access logs for Elastic Load Balancing capture detailed information for all requests made to your load balancer and stores them as log files in the Amazon S3 bucket that you specify. Each log contains details such as the time a request was received, the client's IP address, latencies, request path, and server responses. You can use these access logs to analyze traffic patterns and to troubleshoot your back-end applications. For more information, see [Monitor Your Load Balancer Using Elastic Load Balancing Access Logs](#).

NEW QUESTION 299

If you want to launch Amazon Elastic Compute Cloud (EC2) instances and assign each instance a predetermined private IP address you should:

- A. Launch the instance from a private Amazon Machine Image (AMI).
- B. Assign a group of sequential Elastic IP address to the instances.
- C. Launch the instances in the Amazon Virtual Private Cloud (VPC).
- D. Launch the instances in a Placement Group.
- E. Use standard EC2 instances since each instance gets a private Domain Name Service (DNS) already

Answer: C

Explanation:

Each instance in a VPC has a default network interface (eth0) that is assigned the primary private IP address.

NEW QUESTION 302

You are configuring your company's application to use Auto Scaling and need to move user state information. Which of the following AWS services provides a shared data store with durability and low latency?

- A. AWS ElastiCache Memcached
- B. Amazon Simple Storage Service
- C. Amazon EC2 instance storage
- D. Amazon DynamoDB

Answer: D

Explanation:

https://media.amazonwebservices.com/AWS_Storage_Options.pdf

To speed access to relevant data, many developers pair Amazon S3 with a database, such as Amazon DynamoDB or Amazon RDS. Amazon S3 stores the actual information, and the database serves as the repository for associated metadata (e.g., object name, size, keywords, and so on). Metadata in the database can easily be indexed and queried, making it very efficient to locate an object's reference via a database query. This result can then be used to pinpoint and then retrieve the object itself from Amazon S3.

NEW QUESTION 306

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-Provided 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 307

Which Amazon Elastic Compute Cloud feature can you query from within the instance to access instance properties?

- A. Instance user data
- B. Resource tags
- C. Instance metadata
- D. Amazon Machine Image

Answer: C

Explanation:

Although you can only access instance metadata and user data from within the instance itself, the data is not protected by cryptographic methods.
<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-metadata.html#instancedata-retrieval>

NEW QUESTION 311

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

- A. Security Group Inbound Rule: Protocol - TC

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

Answer: A

NEW QUESTION 314

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

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

Answer: A

Explanation:

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

"Scheduled Reserved

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

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

NEW QUESTION 317

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 322

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

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

Answer: AC

Explanation:

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

NEW QUESTION 324

Which set of Amazon S3 features helps to prevent and recover from accidental data loss?

- A. Object lifecycle and service access logging
- B. Object versioning and Multi-factor authentication
- C. Access controls and server-side encryption
- D. Website hosting and Amazon S3 policies

Answer: B

Explanation:

Versioning-enabled buckets enable you to recover objects from accidental deletion or overwrite. In addition to that, they have made it a requirement that delete operations on versioned data can only be done using MFA (Multi factor authentication).

NEW QUESTION 328

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 333

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 334

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

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

Answer: B

Explanation:

Using multipart upload provides the following advantages:

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

NEW QUESTION 339

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

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

Answer: B

Explanation:

References:

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

NEW QUESTION 340

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 342

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

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

Answer: A

NEW QUESTION 346

A customer has a single 3-TB volume on-premises that is used to hold a large repository of images and print layout files. This repository is growing at 500 GB a year and must be presented as a single logical volume. The customer is becoming increasingly constrained with their local storage capacity and wants an off-site backup of this data, while maintaining low-latency access to their frequently accessed data. Which AWS Storage Gateway configuration meets the customer requirements?

- A. Gateway-Cached volumes with snapshots scheduled to Amazon S3
- B. Gateway-Stored volumes with snapshots scheduled to Amazon S3
- C. Gateway-Virtual Tape Library with snapshots to Amazon S3
- D. Gateway-Virtual Tape Library with snapshots to Amazon Glacier

Answer: A

NEW QUESTION 347

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 348

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

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

Answer: A

Explanation:

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

NEW QUESTION 353

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

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

Answer: AD

NEW QUESTION 354

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

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

Answer: B

Explanation:

Use roles for applications that run on Amazon EC2 instances.

Applications that run on an Amazon EC2 instance need credentials in order to access other AWS services. To provide credentials to the application in a secure way, use IAM roles. A role is an entity that has its own set of permissions, but that isn't a user or group. Roles also don't have their own permanent set of

credentials the way IAM users do. In the case of Amazon EC2, IAM dynamically provides temporary credentials to the EC2 instance, and these credentials are automatically rotated for you.

<http://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html#use-roles-with-ec2>

NEW QUESTION 355

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

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

Answer: C

Explanation:

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

NEW QUESTION 358

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

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

Answer: D

Explanation:

Disabling Source/Destination Checks.

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

NEW QUESTION 362

With which AWS orchestration service can you implement Chef recipes?

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

Answer: C

NEW QUESTION 364

You need to add a route to your routing table in order to allow connections to the internet from your subnet. What route should you add?

- A. Destination: 192.168.1.258/0 --> Target: your Internet gateway
- B. Destination: 0.0.0.0/33 --> Target: your virtual private gateway
- C. Destination: 0.0.0.0/0 --> Target: 0.0.0.0/24
- D. Destination: 10.0.0.0/32 --> Target: your virtual private gateway
- E. Destination: 0.0.0.0/0 --> Target: your Internet gateway

Answer: E

NEW QUESTION 367

You have developed a new web application in us-west-2 that requires six Amazon Elastic Compute Cloud (EC2) instances running at all times. You have three availability zones available in that region (us-west-2a, us-west-2b, and us-west-2c). You need 100 percent fault tolerance if any single Availability Zone in us-west-2 becomes unavailable. How would you do this, each answer has 2 answers, select the answer with BOTH correct answers.

- A. Answer 1 - Us-west-2a with two EC2 instances, us-west-2b with two EC2 instances, and us-west-2c with two EC2 instance
- B. Answer 2 - Us-west-2a with six EC2 instances, us-west-2b with six EC2 instances, and us-west-2c with no EC2 instances
- C. Answer 1 - Us-west-2a with six EC2 instances, us-west-2b with six EC2 instances, and us-west-2c with no EC2 instance
- D. Answer 2 - Us-west-2a with three EC2 instances, us-west-2b with three EC2 instances, and us-west-2c with three EC2 instances.
- E. Answer 1 - Us-west-2a with three EC2 instances, us-west-2b with three EC2 instances, and uswest- 2c with no EC2 instance
- F. Answer 2 - Us-west-2a with three EC2 instances, us-west-2b withthree EC2 instances, and us-west-2c with three EC2 instances.
- G. Answer 1 - Us-west-2a with three EC2 instances, us-west-2b with three EC2 instances, and uswest- 2c with three EC2 instance
- H. Answer 2 - Us-west-2a with four EC2 instances, us-west-2b withtwo EC2 instances, and us-west-2c with two EC2 instance

Answer: B

NEW QUESTION 372

You work for a major news network in Europe. They have just released a new app which allows users to report on events as and when they happen using their

mobile phone. Users are able to upload pictures from the app and then other users will be able to view these pics. Your organization expects this app to grow very quickly, essentially doubling 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 377

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 382

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

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

Answer: A

NEW QUESTION 384

You are a solutions architect who has been asked to do some consulting for a US company that produces re-useable rocket parts. They have a new web application that needs to be built and this application must be stateless. Which three services could you use to achieve this?

- A. AWS Storage Gateway, ElastiCache & ELB
- B. ELB, ElastiCache & RDS
- C. CloudWatch, RDS & DynamoDB
- D. RDS, DynamoDB & ElastiCache

Answer: D

NEW QUESTION 389

You run an automobile reselling company that has a popular online store on AWS. The application sits behind an Auto Scaling group and requires new instances of the Auto Scaling group to identify their public and private IP addresses. How can you achieve this?

- A. By using Ipconfig for windows or Ifconfig for Linux.
- B. By using a cloud watch metric.
- C. Using a Curl or Get Command to get the latest meta-data from <http://169.254.169.254/latest/meta-data/>
- D. Using a Curl or Get Command to get the latest user-data from <http://169.254.169.254/latest/userdata/>

Answer: C

NEW QUESTION 391

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

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

Answer: C

NEW QUESTION 392

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 dashboard. 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 397

You are working in the media industry and you have created a web application where users will be able to upload photos they create to your website. This web application must be able to call the S3 API in order to be able to function. Where should you store your API credentials whilst maintaining the maximum level of security.

- A. Save the API credentials to your php files.
- B. Don't save your API credential
- C. Instead create a role in IAM and assign this role to an EC2 instance when you first create it.
- D. Save your API credentials in a public Github repository.
- E. Pass API credentials to the instance using instance userdat

Answer: B

NEW QUESTION 398

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

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

Answer: A

NEW QUESTION 403

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 406

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 410

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 411

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

- A. True
- B. False

Answer: B

NEW QUESTION 415

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

- A. True
- B. False

Answer: A

NEW QUESTION 419

You can have 1 subnet stretched across multiple availability zones.

- A. True
- B. False

Answer: B

NEW QUESTION 420

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

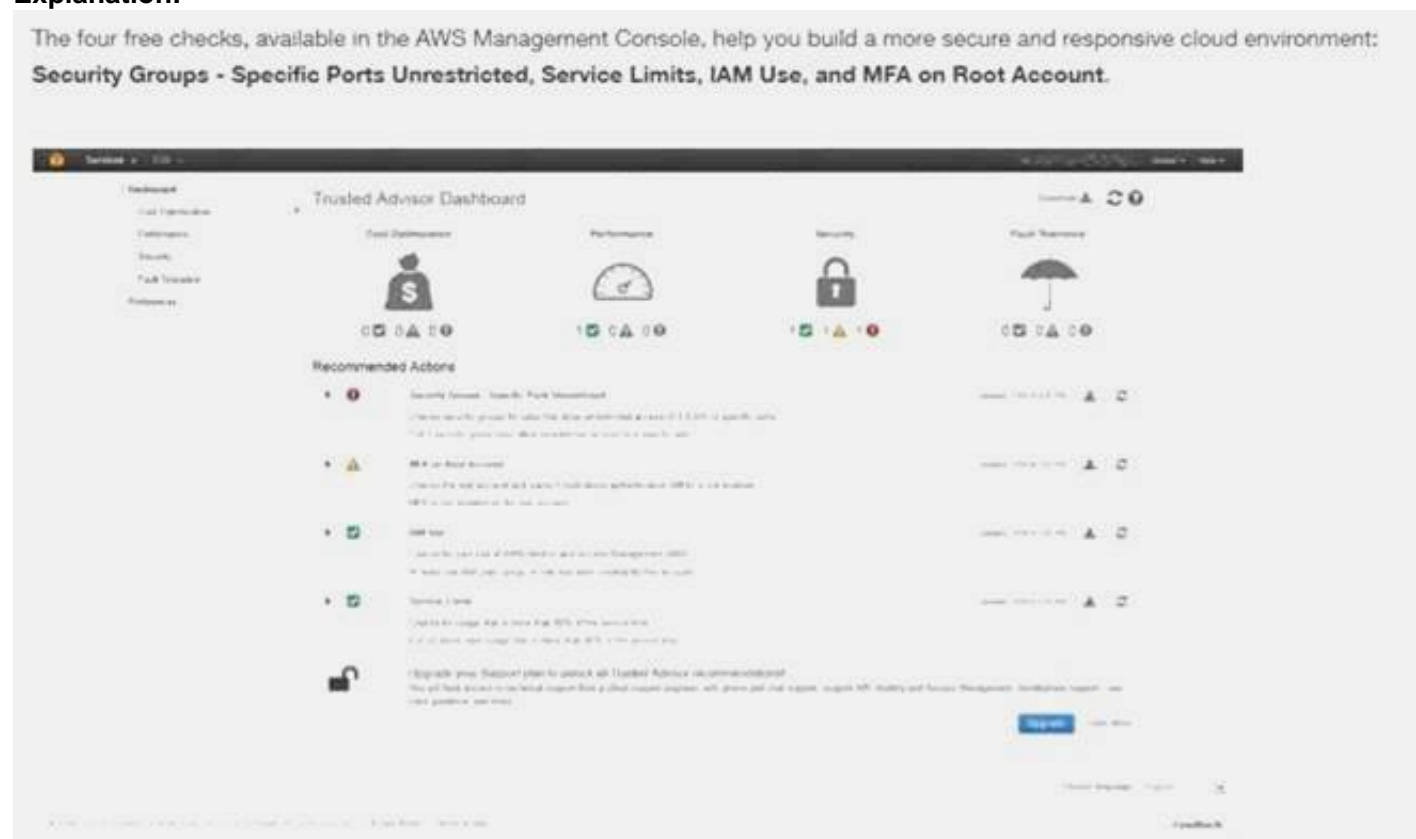
NEW QUESTION 424

Which of the following is not a service of the security category of the AWS trusted advisor service?

- A. Security Groups - Specific Ports Unrestricted
- B. MFA on Root Account
- C. IAM Use
- D. Vulnerability scans on existing VPC

Answer: D

Explanation:



NEW QUESTION 427

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

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

Answer: E

Explanation:

Overview of AWS Import/Export

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

NEW QUESTION 428

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

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

Answer: A

NEW QUESTION 429

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 432

Amazon S3 buckets in all Regions provide which of the following?

- A. Read-after-write consistency for PUTS of new objects AND Strongly consistent for POST & DELETES
- B. Read-after-write consistency for POST of new objects AND Eventually consistent for overwrite PUTS & DELETES
- C. Read-after-write consistency for PUTS of new objects AND Eventually consistent for overwrite PUTS & DELETES
- D. Read-after-write consistency for POST of new objects AND Strongly consistent for POST & DELETES

Answer: C

Explanation:

Q: What data consistency model does Amazon S3 employ?

Amazon S3 buckets in all Regions provide read-after-write consistency for PUTS of new objects and eventual consistency for overwrite PUTS and DELETES.

NEW QUESTION 437

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 442

What is the default VPC security group limit?

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

Answer: A

NEW QUESTION 445

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

NEW QUESTION 448

What is the default per account limit of Elastic IPs?

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

Answer: C

NEW QUESTION 453

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 455

SQL Server stores logins and passwords in the master database.

- A. True
- B. False

Answer: A

Explanation:

There are two authentications

Windows authentication

The credentials for which are not stored in SQL Server database and managed by windows/AD. There would be entry for windows authenticated logins in master database with respective SID but password would be with Active directory.

SQL Server authentication.

For 2nd we have password stored in hash format you can see it from sys.sql_logins. The information about SQL server logins are stored in master database and each login has SID respective to it. Only SA login has same SID no matter what server it is. That is why when you move database by backup restore mechanism users are moved not logins and you finally have to create logins(if already not there) and map it to users. This is generally called as troubleshooting orphaned users

NEW QUESTION 460

Disabling automated backups disables the point-in-time recovery feature.

- A. True
- B. False

Answer: A

NEW QUESTION 461

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

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

Answer: A

Explanation:

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

NEW QUESTION 466

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 467

Amazon Glacier is designed for: (Choose two.)

- A. active database storage.
- B. infrequently accessed data.
- C. data archives.
- D. frequently accessed data.
- E. cached session data

Answer: BC

NEW QUESTION 472

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

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

Answer: B

Explanation:

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

NEW QUESTION 473

A customer's nightly EMR job processes a single 2-TB data file stored on Amazon Simple Storage Service (S3). The EMR job runs on two On-Demand core nodes and three On-Demand task nodes. Which of the following may help reduce the EMR job completion time?
Choose 2 answers

- A. Use three Spot Instances rather than three On-Demand instances for the task nodes.
- B. Change the input split size in the MapReduce job configuration.
- C. Use a bootstrap action to present the S3 bucket as a local filesystem.
- D. Launch the core nodes and task nodes within an Amazon Virtual Cloud.
- E. Adjust the number of simultaneous mapper tasks.
- F. Enable termination protection for the job flo

Answer: BE

NEW QUESTION 474

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 479

What combination of the following options will protect S3 objects from both accidental deletion and accidental overwriting?
Choose 2 answers

- A. Enable S3 versioning on the bucket.
- B. Access S3 data using only signed URLs.
- C. Disable S3 delete using an IAM bucket policy.
- D. Enable S3 Reduced Redundancy Storage.
- E. Enable multi-factor authentication (MFA) protected acces

Answer: AE

NEW QUESTION 484

A VPC public subnet is one that:

- A. Has at least one route in its associated routing table that uses an Internet Gateway (IGW).
- B. Includes a route in its associated routing table via a Network Address Translation (NAT) instance.
- C. Has a Network Access Control List (NACL) permitting outbound traffic to 0.0.0.0/0.
- D. Has the Public Subnet option selected in its configuratio

Answer: A

Explanation:

http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

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

NEW QUESTION 485

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

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

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

Answer: C

NEW QUESTION 494

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 496

Which of the following databases is not supported on Amazon RDS?

- A. MSSOL
- B. MySQL
- C. Aurora
- D. DB2

Answer: D

Explanation:

DB2 is not yet supported on aws. To get the latest list of RDS's supported by AWS, please use the following link :

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

The correct answer is: DB2

NEW QUESTION 498

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 500

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

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

Answer: A

NEW QUESTION 501

You are launching an application in an Auto Scaling group. To store the user session state, you need a structured storage service with durability and low latency. Which service meets your needs?

- A. Amazon ElastiCache
- B. Amazon S3
- C. Amazon EC2 instance storage
- D. Amazon DynamoDB

Answer: D

NEW QUESTION 502

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

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

Answer: A

NEW QUESTION 503

Your company has set up an application in eu-west-1 with a disaster recovery site in eu-central-1. You want to be notified of any AWS API activity in regions other than these two. How can you monitor AWS API activity in other regions?

- A. Create a CloudWatch alarm for CloudTrail events.
- B. Create a CloudWatch alarm for Trusted Advisor.
- C. Create a CloudWatch alarm for VPC flow logs.
- D. Create a CloudWatch alarm for SSH key usage.

Answer: A

NEW QUESTION 505

You have a Cassandra cluster running in private subnets in an Amazon VPC. A new application in a different Amazon VPC needs access to the database. How can the new application access the database?

- A. Set up a dual-homed instance with ENIs in both Amazon VPCs.
- B. Set up a VPC peering connection between the two Amazon VPCs.
- C. Set up a NAT Gateway in the database's Amazon VPC.
- D. Set up a NAT Gateway in the application's Amazon VPC.

Answer: C

NEW QUESTION 507

You bid \$0.22 for an Amazon EC2 Spot Instance when the market price was \$0.20. For 90 minutes, the market price remained at \$0.20. Then the market price

changed to \$0.25, and your instance was terminated by AWS.
What was your cost of running the instance for the entire duration?

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

Answer: D

NEW QUESTION 510

Your organization is looking for a solution that can help the business with streaming data. Several services will require access to read and process the same stream concurrently. What AWS service meets the business requirements?

- A. Amazon Kinesis Firehose
- B. Amazon Kinesis Streams
- C. Amazon CloudFront
- D. Amazon SQS

Answer: B

NEW QUESTION 511

A customer's security team requires the logging of all network access attempts to Amazon EC2 instances in their production VPC on AWS.
Which configuration will meet the security team's requirement?

- A. Enable CloudTrail for the production VPC.
- B. Enable VPC Flow Logs for the production VPC.
- C. Enable both CloudTrail and VPC Flow Logs for the production VPC.
- D. Enable both CloudTrail and VPC Flow Logs for the AWS account.

Answer: B

NEW QUESTION 512

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual SAA-C01 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the SAA-C01 Product From:

<https://www.2passeasy.com/dumps/SAA-C01/>

Money Back Guarantee

SAA-C01 Practice Exam Features:

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