

SAA-C01 Dumps

AWS Certified Solutions Architect - Associate

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



NEW QUESTION 1

A 3-tier e-commerce web application is currently deployed on-premises and will be migrated to AWS for greater scalability and elasticity. The web server currently shares read-only data using a network distributed file system. The app server tier uses a clustering mechanism for discovery and shared session state that depends on IP multicast. The database tier uses shared-storage clustering to provide database 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 time
- B. App servers: share state using a combination of DynamoDB and IP unicast
- C. Database: use RDS with multi-AZ deployment and one or more read replicas
- D. Backup: web servers, app servers, and database backed up weekly to Glacier using snapshots.
- E. Web servers: store read-only data in an EC2 NFS server, mount to each web server at boot time
- F. App servers: share state using a combination of DynamoDB and IP multicast
- G. Database: use RDS with multi-AZ deployment and one or more Read Replicas
- H. Backup: web and app servers backed up weekly via AMIs, database backed up via DB snapshots.
- I. Web servers: store read-only data in S3, and copy from S3 to root volume at boot time
- J. App servers: share state using a combination of DynamoDB and IP unicast
- K. Database: use RDS with multi-AZ deployment and one or more Read Replicas
- L. Backup: web and app servers backed up weekly via AMIs, database backed up via DB snapshots.
- M. Web servers: store read-only data in S3, and copy from S3 to root volume at boot time
- N. App servers: share state using a combination of DynamoDB and IP unicast
- O. Database: use RDS with multi-AZ deployment
- P. Backup: web and app servers backed up weekly via AMIs, database backed up via DB snapshots.

Answer: 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

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

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

Answer: B

NEW QUESTION 3

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 4

You have launched an EC2 instance with four (4) 500 GB EBS Provisioned IOPS volumes attached. The EC2 Instance is EBS-Optimized and supports 500 Mbps throughput between EC2 and EBS. The two EBS volumes are configured as a single RAID 0 device, and each Provisioned IOPS volume is provisioned with 4,000 IOPS (4,000 16KB reads or writes) for a total of 16,000 random IOPS on the instance. The EC2 Instance initially delivers the expected 16,000 IOPS random read and write performance. Sometime later in order to increase the total random I/O performance of the instance, you add an additional two 500 GB EBS Provisioned IOPS volumes to the RAID. Each volume is provisioned to 4,000 IOPS like the original four for a total of 24,000 IOPS on the EC2 instance. Monitoring shows that the EC2 instance CPU utilization increased from 50% to 70%. but the total random IOPS measured at the instance level does not increase at all.

What is the problem and a valid solution?

- A. Larger storage volumes support higher Provisioned IOPS rates: increase the provisioned volume storage of each of the 6 EBS volumes to 1TB
- B. The EBS-Optimized throughput limits the total IOPS that can be utilized: use an EBS-Optimized instance that provides larger throughput.
- C. Small block sizes cause performance degradation, limiting the I/O throughput, configure the instance device driver and file system to use 64KB blocks to increase throughput.
- D. RAID 0 only scales linearly to about 4 devices, use RAID 0 with 4 EBS Provisioned IOPS volumes but increase each Provisioned IOPS EBS volume to 6,000 IOPS.
- E. The standard EBS instance root volume limits the total IOPS rate, change the instance root volume to also be a 500GB 4,000 Provisioned IOPS volume.

Answer: E

NEW QUESTION 5

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

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

Answer: B

Explanation:

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

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

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

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

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

NEW QUESTION 6

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

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

Answer: AB

Explanation:

The Route 53 and IAM are global.

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

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

NEW QUESTION 7

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

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

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

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

Answer: A

Explanation:

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

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

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

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

NEW QUESTION 8

You are responsible for a legacy web application whose server environment is approaching end of life. You would like to migrate this application to AWS as quickly

as possible, since the application environment currently has the following limitations:

The VM's single 10GB VMDK is almost full

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

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 9

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

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

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

Answer: BE

NEW QUESTION 10

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 point
- 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 10

You are designing the network infrastructure for an application server in Amazon VPC. Users will

access all the application instances from the Internet as well as from an on-premises network. The on-premises network is connected to your VPC over an AWS Direct Connect link.

How would you design routing to meet the above requirements?

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

Answer: B

NEW QUESTION 13

You are implementing AWS Direct Connect. You intend to use AWS public service end points such as Amazon S3, across the AWS Direct Connect link. You want other Internet traffic to use your existing link to an Internet Service Provider.

What is the correct way to configure AWS Direct connect for access to services such as Amazon S3?

- A. Configure a public Interface on your AWS Direct Connect link
- B. Configure a static route via your AWS Direct Connect link that points to Amazon S3 Advertise a default route to AWS using BGP.
- C. Create a private interface on your AWS Direct Connect link
- D. Configure a static route via your AWS Direct connect link that points to Amazon S3 Configure specific routes to your network in your VPC,
- E. Create a public interface on your AWS Direct Connect link

- F. Redistribute BGP routes into your existing routing infrastructure advertise specific routes for your network to AWS.
- G. Create a private interface on your AWS Direct connect lin
- H. Redistribute BGP routes into your existing routing infrastructure and advertise a default route to AWS.

Answer: C

Explanation:

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

NEW QUESTION 15

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

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

Answer: D

Explanation:

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

ELB supports sticky sessions only in HTTP/HTTPS configurations.

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

terminated on the back-end instance, the Elastic Load Balancer has no visibility into the HTTP headers and therefore cannot set or read any of the HTTP headers being passed back and forth. <http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/elb-sticky-sessions.html>

NEW QUESTION 19

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

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

Answer: D

Explanation:

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

NEW QUESTION 24

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

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

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

Answer: D

Explanation:

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

Q: What are some EMR best practices?

If you are running EMR in production you should specify an AMI version, Hive version, Pig version, etc. to make sure the version does not change unexpectedly (e.g. when EMR later adds support for a newer version). If your cluster is mission critical, only use Spot instances for task nodes because if the Spot price increases you may lose the instances. In development, use logging and enable debugging

to spot and correct errors faster. If you are using GZIP, keep your file size to 1–2 GB because GZIP files cannot be split. Click here to download the white paper on Amazon EMR best practices. <https://aws.amazon.com/elasticmapreduce/faqs>

NEW QUESTION 28

You currently operate a web application. In the AWS US-East region The application runs on an autoscaled layer of EC2 instances and an RDS Multi-AZ database

Your IT security compliance officer has tasked you to develop a reliable and durable logging solution to track changes made to your EC2.IAM And RDS resources. The solution must ensure the integrity and confidentiality of your log data. Which of these solutions would you recommend?

- A. Create a new CloudTrail trail with one new S3 bucket to store the logs and with the global services option selecte
- B. Use IAM roles S3 bucket policies and Multi Factor Authentication (MFA). Delete on the S3 bucket that stores your logs.
- C. Create a new CloudTrail with one new S3 bucket to store the log
- D. Configure SNS to send log file delivery notifications to your management syste
- E. Use IAM roles and S3 bucket policies on the S3 bucket mat stores your logs.
- F. Create a new CloudTrail trail with an existing S3 bucket to store the logs and with the global services option selecte
- G. Use S3 ACLs and Multi Factor Authentication (MFA). Delete on the S3 bucket that stores your logs.
- H. Create three new CloudTrail trails with three new S3 buckets to store the logs one for the AWS Management console, one for AWS SDKs and one for command line tool
- I. Use IAM roles and S3 bucket policies on the S3 buckets that store your logs.

Answer: A

NEW QUESTION 30

An enterprise wants to use a third-party SaaS application. The SaaS application needs to have access to issue several API commands to discover Amazon EC2 resources running within the enterprise's account The enterprise has internal security policies that require any outside access to their environment must conform to the principles of least privilege and there must be controls in place to ensure that the credentials used by the SaaS vendor cannot be used by any other third party. Which of the following would meet all of these conditions?

- A. From the AWS Management Console, navigate to the Security Credentials page and retrieve the access and secret key for your account.
- B. Create an IAM user within the enterprise account assign a user policy to the IAM user that allows only the actions required by the SaaS application create a new access and secret key for the user andprovide these credentials to the SaaS provider.
- C. Create an IAM role for cross-account access allows the SaaS provider's account to assume the role and assign it a policy that allows only the actions required by the SaaS application.
- D. Create an IAM role for EC2 instances, assign it a policy that allows only the actions required for the Saas application to work, provide the role ARM to the SaaS provider to use when launching their application instances.

Answer: C

Explanation:

Granting Cross-account Permission to objects It Does Not Own

In this example scenario, you own a bucket and you have enabled other AWS accounts to upload objects. That is, your bucket can have objects that other AWS accounts own.

Now, suppose as a bucket owner, you need to grant cross-account permission on objects, regardless of who the owner is, to a user in another account. For example, that user could be a billing application that needs to access object metadata. There are two core issues:

The bucket owner has no permissions on those objects created by other AWS accounts. So for the bucket owner to grant permissions on objects it does not own, the object owner, the AWS account that created the objects, must first grant permission to the bucket owner. The bucket owner can then delegate those permissions.

Bucket owner account can delegate permissions to users in its own account but it cannot delegate permissions to other AWS accounts, because cross-account delegation is not supported.

In this scenario, the bucket owner can create an AWS Identity and Access Management (IAM) role with permission to access objects, and grant another AWS account permission to assume the role temporarily enabling it to access objects in the bucket.

Background: Cross-Account Permissions and Using IAM Roles

IAM roles enable several scenarios to delegate access to your resources, and cross-account access is one of the key scenarios. In this example, the bucket owner, Account A, uses an IAM role to temporarily delegate object access cross-account to users in another AWS account, Account C. Each IAM role you create has two policies attached to it:


A trust policy identifying another AWS account that can assume the role.

An access policy defining what permissions—for example, s3:GetObject—are allowed when someone assumes the role. For a list of permissions you can specify in a policy, see Specifying Permissions in a Policy.

The AWS account identified in the trust policy then grants its user permission to assume the role. The user can then do the following to access objects:

Assume the role and, in response, get temporary security credentials. Using the temporary security credentials, access the objects in the bucket.

For more information about IAM roles, go to Roles (Delegation and Federation) in IAM User Guide. The following is a summary of the walkthrough steps:



Account A administrator user attaches a bucket policy granting Account B conditional permission to upload objects.

Account A administrator creates an IAM role, establishing trust with Account C, so users in that account can access Account A. The access policy attached to the role limits what user in Account C can do when the user accesses Account A.

Account B administrator uploads an object to the bucket owned by Account A, granting full-control permission to the bucket owner.

Account C administrator creates a user and attaches a user policy that allows the user to assume the role.

User in Account C first assumes the role, which returns the user temporary security credentials. Using those temporary credentials, the user then accesses objects in the bucket.

For this example, you need three accounts. The following table shows how we refer to these accounts and the administrator users in these accounts. Per IAM

guidelines (see About Using an Administrator User to Create Resources and Grant Permissions) we do not use the account root credentials in this walkthrough. Instead, you create an administrator user in each account and use those credentials in creating resources and granting them permissions

AWS Account ID	Account Referred To As	Administrator User in the Account
1111-1111-1111	Account A	AccountAdmin
2222-2222-2222	Account B	AccountBadmin
3333-3333-3333	Account C	AccountCadmin

NEW QUESTION 35

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 38

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 39

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 text 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 parallel
- 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 elaboration commands to a group of hosts working in parallel
- 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 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 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 43

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 capturing 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 48

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 51

A company is running a batch analysis every hour on their main transactional DB. running on an RDS MySQL instance to populate their central Data Warehouse running on Redshift During the execution of the batch their transactional applications are very slow When the batch completes they need to update the top management dashboard with the new data The dashboard is produced by another system running on-premises that is currently started when a manually-sent email notifies that an update is required The on-premises system cannot be modified because is managed by another team. How would you optimize this scenario to solve performance issues and automate the process as much as possible?

- A. Replace RDS with Redshift for the batch analysis and SNS to notify the on-premises system to update the dashboard
- B. Replace ROS with Redsnift for the oaten analysis and SQS to send a message to the on-premises system to update the dashboard
- C. Create an RDS Read Replica for the batch analysis and SNS to notify me on-premises system to update the dashboard
- D. Create an RDS Read Replica for the batch analysis and SQS to send a message to the on-premises system to update the dashboard.

Answer: C

Explanation:

If you want to prevent your reporting and analytic processing from interfering with the performance of your OLTP workload."

If I understand the above statement correctly, they are saying to separate reporting and analytic processing from OLTP. In other word, use RedShift for reporting and analytic processing and use RDS for OLTP workload.

NEW QUESTION 56

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

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

Answer: AB

NEW QUESTION 59

To serve Web traffic for a popular product your chief financial officer and IT director have purchased 10 ml large heavy utilization Reserved Instances (RIs) evenly spread across two availability zones: Route 53 is used to deliver the traffic to an Elastic Load Balancer (ELB). After several months, the product grows even more popular and you need additional capacity. As a result, your company purchases two C3.2xlarge medium utilization Ris. You register the two c3 2xlarge instances with your ELB and quickly find that the ml large instances are at 100% of capacity and the c3 2xlarge instances have significant capacity that's unused. Which option is the most cost effective and uses EC2 capacity most effectively?

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

Answer: A

Explanation:

Weighted Routing Policy

Use the weighted routing policy when you have multiple resources that perform the same function (for example, web servers that serve the same website) and you

want Amazon Route 53 to route traffic to those resources in proportions that you specify (for example, one quarter to one server and three quarters to the other). For more information about weighted resource record sets, see [Weighted Routing](#).

NEW QUESTION 62

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

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

Answer: A

Explanation:

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

NEW QUESTION 67

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

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

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

Answer: BCE

Explanation:

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

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

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

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

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

NEW QUESTION 70

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

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

Answer: A

Explanation:

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

NEW QUESTION 71

Your team has a tomcat-based Java application you need to deploy into development, test and production environments. After some research, you opt to use Elastic Beanstalk due to its tight integration with your developer tools and RDS due to its ease of management. Your QA team lead points out that you need to roll a sanitized set of production data into your environment on a nightly basis. Similarly, other software teams in your org want access to that same restored data via their EC2 instances in your VPC. The optimal setup for persistence and security that meets the above requirements would be the following.

- A. Create your RDS instance as part of your Elastic Beanstalk definition and alter its security group to allow access to it from hosts in your application subnets.
- B. Create your RDS instance separately and add its IP address to your application's DB connection strings in your code Alter its security group to allow access to it from hosts within your VPC's IP address block.
- C. Create your RDS instance separately and pass its DNS name to your app's DB connection string as an environment variabl

- D. Create a security group for client machines and add it as a valid source for DB traffic to the security group of the RDS instance itself.
- E. Create your RDS instance separately and pass its DNS name to your's DB connection string as an environment variable Alter its security group to allow access to it from hosts in your application subnets.

Answer: C

Explanation:

Elastic Beanstalk provides support for running Amazon RDS instances in your Elastic Beanstalk environment. This works great for development and testing environments, but is not ideal for a production environment because it ties the lifecycle of the database instance to the lifecycle of your application's environment. It can't be D because RDS is opened to all "hosts in your application subnets" where C only opens RDS to specific client machines in a specific security group. <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/AWSHowTo.RDS.html>

NEW QUESTION 73

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 Billin
- G. This will give IAM users in the Master account access to resources in the Dev and Test accounts

Answer: C

NEW QUESTION 75

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 76

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

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

Answer: B

Explanation:

You only need one stack to contain two layers:

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

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

NEW QUESTION 80

What does Amazon S3 stand for?

- A. Simple Storage Solution.
- B. Storage Storage Storage (triple redundancy Storage).
- C. Storage Server Solution.
- D. Simple Storage Servic

Answer: D

Explanation:

Amazon Simple Storage Service (Amazon S3) is storage for the Internet. It provides a simple interface to manage scalable, reliable, and low latency data storage service over the Internet. <http://docs.aws.amazon.com/AmazonS3/latest/gsg/GetStartedWithS3.html>

NEW QUESTION 81

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 82

Select the most correct answer

The device name /dev/sda1 (within Amazon EC2) is _____

- A. Possible for EBS volumes
- B. Reserved for the root device
- C. Recommended for EBS volumes
- D. Recommended for instance store volumes

Answer: B

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/RootDeviceStorage.html> The root device is typically /dev/sda1 (Linux) or xvda (Windows).

NEW QUESTION 85

If I want an instance to have a public IP address, which IP address should I use?

- A. Elastic IP Address
- B. Class B IP Address
- C. Class A IP Address
- D. Dynamic IP Address

Answer: A

NEW QUESTION 90

What does Amazon SWF stand for?

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

Answer: B

NEW QUESTION 91

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 94

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 99

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

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

Answer: A

NEW QUESTION 103

Every user you create in the IAM system starts with .

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

Answer: C

NEW QUESTION 104

Can you create IAM security credentials for existing users?

- A. Yes, existing users can have security credentials associated with their account.
- B. No, IAM requires that all users who have credentials set up are not existing users
- C. No, security credentials are created within GROUPS, and then users are associated to GROUPS at a later time.
- D. Yes, but only IAM credentials, not ordinary security credential

Answer: A

NEW QUESTION 106

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

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

Answer: D

Explanation:

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

NEW QUESTION 109

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

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

Answer: C

NEW QUESTION 113

What does Amazon Elastic Beanstalk provide?

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

Answer: B

NEW QUESTION 117

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 122

What will be the status of the snapshot until the snapshot is complete.

- A. running
- B. working
- C. progressing
- D. pending

Answer: D

Explanation:

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

Creating an Amazon EBS Snapshot

After writing data to an EBS volume, you can periodically create a snapshot of the volume to use as a baseline for new volumes or for data backup. If you make periodic snapshots of a volume, the snapshots are incremental so that only the blocks on the device that have changed after your last snapshot are saved in the new snapshot. Even though snapshots are saved incrementally, the snapshot deletion process is designed so that you need to retain only the most recent snapshot in order to restore the volume.

Snapshots occur asynchronously; the point-in-time snapshot is created immediately, but the status of the snapshot is `pending` until the snapshot is complete (when all of the modified blocks have been transferred to Amazon S3), which can take several hours for large initial snapshots or subsequent snapshots where many blocks have changed. While it is completing, an in-progress snapshot is not affected by ongoing reads and writes to the volume.

NEW QUESTION 123

True or False: Automated backups are enabled by default for a new DB Instance.

- A. TRUE
- B. FALSE

Answer: A

NEW QUESTION 125

What does the AWS Storage Gateway provide?

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

Answer: A

NEW QUESTION 128

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

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

Answer: D

Explanation:

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

NEW QUESTION 133

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 136

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 141

Amazon RDS DB snapshots and automated backups are stored in

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

Answer: A

NEW QUESTION 143

Groups can't _____.

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

Answer: B

Explanation:

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

NEW QUESTION 148

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

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

Answer: D

NEW QUESTION 149

Changes to the backup window take effect _____.

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

Answer: C

Explanation:

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

NEW QUESTION 151

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 154

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 159

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

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

Answer: C

NEW QUESTION 160

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

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

Answer: C

NEW QUESTION 162

What is the maximum write throughput I can provision for a single Dynamic DB table?

- A. 1,000 write capacity units
- B. 100,000 write capacity units
- C. Dynamic DB is designed to scale without limits, but if you go beyond 10,000 you have to contact AWS first.
- D. 10,000 write capacity units

Answer: C

Explanation:

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

Q: Is there a **limit** to how much throughput I can get out of a single table?

No, you can increase the throughput you have provisioned for your table using UpdateTable API or in the AWS Management Console. DynamoDB is able to operate at massive scale and there is no theoretical limit on the maximum throughput you can achieve. DynamoDB automatically divides your table across multiple partitions, where each partition is an independent parallel computation unit. DynamoDB can achieve increasingly high throughput rates by adding more partitions.

If you wish to exceed throughput rates of 10,000 writes/second or 10,000 reads/second, you must first contact Amazon through this [online form](#).

NEW QUESTION 166

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 170

Can I move a Reserved Instance from one Region to another?

- A. No
- B. Only if they are moving into GovCloud
- C. Yes
- D. Only if they are moving to US East from another region

Answer: A

NEW QUESTION 175

What is the durability of S3 RRS?

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

Answer: A

Explanation:

RRS = Reduced Redundancy Storage

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

NEW QUESTION 177

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

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

Answer: D

Explanation:

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

NEW QUESTION 178

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

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

Answer: B

NEW QUESTION 180

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

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

Answer: A

Explanation:

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

This never states "groups a user can be in"

NEW QUESTION 183

True or False: When you perform a restore operation to a point in time or from a DB Snapshot, a new DB Instance is created with a new endpoint.

- A. FALSE
- B. TRUE

Answer: B

Explanation:

Restoring From a DB Snapshot

Amazon RDS creates a storage volume snapshot of your DB instance, backing up the entire DB instance and not just individual databases. You can create a DB instance by restoring from this DB snapshot. When you restore the DB instance, you provide the name of the DB snapshot to restore from, and then provide a name for the new DB instance that is created from the restore. You cannot restore from a DB snapshot to an existing DB instance; a new DB instance is created when you restore. http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_RestoreFromSnapshot.html

NEW QUESTION 186

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

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

Answer: D

NEW QUESTION 190

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

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

Answer: D

Explanation:

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

NEW QUESTION 191

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 194

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

- A. the checks may still be in progress on the volume
- B. the check has passed
- C. the check has failed

Answer: A

Explanation:

If the status is insufficient-data, the checks may still be in progress on the volume. You can view the results of volume status checks to identify any impaired volumes and take any necessary actions.

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

NEW QUESTION 198

By default, when an EBS volume is attached to a Windows instance, it may show up as any drive letter on the instance. You can change the settings of the _____ Service to set the drive letters of the EBS volumes per your specifications.

- A. EBSConfig Service
- B. AMIConfig Service
- C. Ec2Config Service
- D. Ec2-AMIConfig Service

Answer: C

Explanation:

Ec2Config Service is like sysprep and used specifically for windows instances. You can change parameters in OS before launching.

NEW QUESTION 200


What is Oracle SQL Developer?

- A. An AWS developer who is an expert in Amazon RDS using both the Oracle and SQL Server DB engines
- B. A graphical Java tool distributed without cost by Oracle.
- C. It is a variant of the SQL Server Management Studio designed by Microsoft to support Oracle DBMS functionalities
- D. A different DBMS released by Microsoft free of cost

Answer: B

Explanation:

<http://www.oracle.com/technetwork/developer-tools/sql-developer/what-is-sqldev-093866.html>



NEW QUESTION 201

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

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

Answer: C

Explanation:

An IAM group is a collection of IAM users. You can use groups to specify permissions for a collection of users, which can make those permissions easier to

manage for those users. <http://docs.aws.amazon.com/IAM/latest/UserGuide/id.html>

NEW QUESTION 205

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

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

Answer: D

NEW QUESTION 210

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 214

Read Replicas require a transactional storage engine and are only supported for the _____ storage engine

- A. OracleSAM
- B. MSSQLDB
- C. InnoDB
- D. MyISAM

Answer: C

Explanation:

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

- Using a non-transactional storage engine such as MyISAM. Read replicas require a transactional storage engine. Replication is only supported for the **InnoDB** storage engine on MySQL and the XtraDB storage engine on MariaDB.

NEW QUESTION 215

When should I choose Provisioned IOPS over Standard RDS storage?

- A. If you have batch-oriented workloads
- B. If you use production online transaction processing (OLTP) workloads.
- C. If you have workloads that are not sensitive to consistent performance

Answer: B

Explanation:

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

Amazon RDS provisions that IOPS rate and storage for the lifetime of the DB instance or until you change it. Provisioned IOPS storage is optimized for I/O intensive, online transaction processing (OLTP) workloads that have consistent performance requirements. Provisioned IOPS helps performance tuning.

NEW QUESTION 216

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

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

Answer: B

Explanation:

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

NEW QUESTION 218

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 221

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

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

Answer: C

NEW QUESTION 222

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 225

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

- A. No
- B. Yes

Answer: B

NEW QUESTION 227

In regards to IAM you can edit user properties later, but you cannot use the console to change the _____.

- A. user name
- B. password
- C. default group

Answer: A

NEW QUESTION 229

Can I test my DB Instance against a new version before upgrading?

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

Answer: B

Explanation:

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

NEW QUESTION 231

Can I use Provisioned IOPS with VPC?

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

Answer: D

NEW QUESTION 234

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 235

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

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

Answer: A

Explanation:

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

NEW QUESTION 238

What are the four levels of AWS Premium Support?

- A. Basic, Developer, Business, Enterprise
- B. Basic, Startup, Business, Enterprise
- C. Free, Bronze, Silver, Gold
- D. All support is free

Answer: A

Explanation:

Q: How are the enhanced AWS Support tiers different from Basic Support? AWS Basic Support offers all AWS customers access to our Resource Center, Service Health Dashboard, Product FAQs, Discussion Forums, and Support for Health Checks at no additional charge. Customers who desire a deeper level of support can subscribe to AWS Support at the Developer, Business, or Enterprise level. <https://aws.amazon.com/premiumsupport/faqs/>

NEW QUESTION 241

Please select the Amazon EC2 resource which cannot be tagged.

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

Answer: C

Explanation:

http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html#tag-restrictions

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

NEW QUESTION 243

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 248

Through which of the following interfaces is AWS Identity and Access Management available?

- A) AWS Management Console
- B) Command line interface (CLI)
- C) IAM Query API
- D) Existing libraries

- A. Only through Command line interface (CLI)
- B. A, B and C
- C. A and C
- D. All of the above

Answer: D

Explanation:

Accessing IAM:

1 - AWS Management Console 2 - AWS Command Line Tools

3 - AWS SDKs (i.e. Existing libraries) 4 - IAM HTTPS API

<http://docs.aws.amazon.com/IAM/latest/UserGuide/introduction.html#intro-accessing>

NEW QUESTION 253

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 258

How are the EBS snapshots saved on Amazon S3?

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

Answer: B

NEW QUESTION 260

What happens when you create a topic on Amazon SNS?

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

Answer: B

NEW QUESTION 264

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 266

The ____ service is targeted at organizations with multiple users or systems that use AWS products such as Amazon EC2, Amazon SimpleDB, and the AWS Management Console.

- A. Amazon RDS
- B. AWS Integrity Management
- C. AWS Identity and Access Management
- D. Amazon EMR

Answer: C

Explanation:

https://aws.amazon.com/documentation/iam/?nc1=h_ls

NEW QUESTION 269

When you use the AWS Management Console to delete an IAM user, IAM also deletes any signing certificates and any access keys belonging to the user.

- A. FALSE
- B. This is configurable
- C. TRUE

Answer: C

Explanation:

When you use the AWS Management Console to delete an IAM user, IAM automatically deletes the following information for you:

The user

Any group memberships -- that is, the user is removed from any IAM groups that the user was a member of:

Any password associated with the user Any access keys belonging to the user

All inline policies embedded in the user (policies that are applied to a user via group permissions are not affected) Note!

Any managed policies attached to the user are detached from the user when the user is deleted. Managed policies are not deleted when you delete a user.

Any associated MFA device http://docs.aws.amazon.com/IAM/latest/UserGuide/id_users_manage.html#id_users_deleting_console

NEW QUESTION 272

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 273

What is the default maximum number of MFA devices in use per AWS account (at the root account level)?

- A. 1
- B. 5
- C. 15
- D. 10

Answer: A

Explanation:

http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_iam-limits.html

NEW QUESTION 274

What does Amazon Route53 provide?

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

Answer: C

Explanation:

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

NEW QUESTION 277

What does Amazon ElastiCache provide?

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

Answer: C

NEW QUESTION 278

How many Elastic IP by default in Amazon Account?

- A. 1 Elastic IP
- B. 3 Elastic IP
- C. 5 Elastic IP
- D. 0 Elastic IP

Answer: C

Explanation:

"By default, all AWS accounts are limited to 5 Elastic IP addresses, because public (IPv4) Internet addresses are a scarce public resource."

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

NEW QUESTION 279

The one-time payment for Reserved Instances is ____ refundable if the reservation is cancelled.

- A. always
- B. in some circumstances
- C. never

Answer: C

Explanation:

the one-time fee is non-refundable.

<https://aws.amazon.com/ec2/purchasing-options/reserved-instances/buyer/>

Important Notes about Purchases

- If your needs change, you can modify or exchange reserved instances, or list eligible Standard Reserved Instances for sale on the Reserved Instance Marketplace.
- You can purchase up to 20 Reserved Instances per Availability Zone each month. If you need additional Reserved Instances, complete the form found [here](#).
- Purchases of Reserved Instances are **non-refundable**.
- If you purchase a Reserved Instance from a third-party seller, we will share your city, state, and zip code with the seller for tax purposes. If you don't wish to purchase from a 3rd party seller, please make sure to select a Reserved Instance with "AWS" listed as the seller in the console purchasing screen.

NEW QUESTION 281

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 283

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 286

A group can contain many users. Can a user belong to multiple groups?

- A. Yes always
- B. No
- C. Yes but only if they are using two factor authentication
- D. Yes but only in VPC

Answer: A

Explanation:

A group can contain many users, and a user can belong to multiple groups. http://docs.aws.amazon.com/IAM/latest/UserGuide/id_groups.html

NEW QUESTION 288

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 292

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

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

Answer: A

Explanation:

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

Amazon EC2 Instance Store

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

NEW QUESTION 297

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

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

Answer: C

NEW QUESTION 300

While controlling access to Amazon EC2 resources, which of the following acts as a firewall that controls the traffic allowed to reach one or more instances?

- A. A security group
- B. An instance type
- C. A storage cluster
- D. An object

Answer: A

Explanation:

A security group acts as a firewall that controls the traffic allowed to reach one or more instances. When you launch an instance, you assign it one or more security groups. <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/UsingIAM.html>

NEW QUESTION 305

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

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

Answer: A

Explanation:

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

NEW QUESTION 306

Which DNS name can only be resolved within Amazon EC2?

- A. Internal DNS name
- B. External DNS name
- C. Global DNS name
- D. Private DNS name

Answer: D

Explanation:

To view DNS hostnames for an instance using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, choose Instances.
3. Select your instance from the list.
4. In the details pane, the Public DNS (IPv4) and Private DNS fields display the DNS hostnames, if applicable.

NEW QUESTION 311

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 316

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 319

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 320

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 323

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

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

Answer: C

Explanation:

Private MFA does not apply here.

Q. What is AWS MFA?

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

NEW QUESTION 326

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

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

Answer: D

NEW QUESTION 329

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 334


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

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

Answer: A

Explanation:

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



NEW QUESTION 338

Location of Instances is _____

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

Answer: B

Explanation:

Regions and Availability Zones

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

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

NEW QUESTION 339

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

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

Answer: B

NEW QUESTION 344

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

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

Answer: C

Explanation:

Amazon EC2 uses SSH keys, Windows passwords, and security groups to control who has access to the operating system of specific Amazon EC2 instances. There's no method in the IAM system to allow or deny access to the operating system of a specific instance.
http://docs.aws.amazon.com/IAM/latest/UserGuide/IAM_UseCases.html

NEW QUESTION 348

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 351

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 353

Within the IAM service a GROUP is regarded as a:

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

Answer: D

Explanation:

Use groups to assign permissions to IAM users

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

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

NEW QUESTION 355

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

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

Answer: B

Explanation:

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

NEW QUESTION 357

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 360

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 361

Which choice is a storage option supported by Amazon EC2?

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

Answer: B

Explanation:

Amazon EC2 supports the following storage options: Amazon Elastic Block Store (Amazon EBS) Amazon EC2 Instance Store Amazon Simple Storage Service (Amazon S3) <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/Storage.html>

NEW QUESTION 362

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

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 364

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

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

Answer: C

Explanation:

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

NEW QUESTION 369

Without _____, you must either create multiple AWS accounts-each with its own billing and subscriptions to AWS products-or your employees must share the security credentials of a single AWS account.

- A. Amazon RDS
- B. Amazon Glacier
- C. Amazon EMR
- D. Amazon IAM

Answer: D

NEW QUESTION 371

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 373

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

- A. private key
- B. foreign key
- C. public key
- D. protected key

Answer: C

Explanation:

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

NEW QUESTION 378

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

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

Answer: A

Explanation:

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

NEW QUESTION 383

In the Amazon RDS which uses the SQL Server engine, what is the maximum size for a Microsoft SQL Server DB Instance with SQL Server Express edition?

- A. 10 GB per DB
- B. 100 GB per DB
- C. 2 TB per DB
- D. 1TB per DB

Answer: A

Explanation:

The maximum storage size for a Microsoft SQL Server DB Instance is 4 TB for all instances except the SQL Server Express edition, which limits storage to a total of 300 GB. The minimum storage size for a Microsoft SQL Server DB Instance is 20 GB for the Microsoft SQL Server Express and Web Editions and 200 GB for the Standard and Enterprise Editions. http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_SQLServer.html

NEW QUESTION 385

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

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

Answer: A

Explanation:

Best Practices for Configuring Elastic Network Interfaces

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

NEW QUESTION 389

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

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

Answer: B

Explanation:

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

NEW QUESTION 394

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 398

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

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

Answer: D

Explanation:

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

NEW QUESTION 401

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

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

F. Put the database behind an Elastic Load Balance

Answer: AE


NEW QUESTION 404

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

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

Answer: A

Explanation:



NEW QUESTION 409

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 413

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 417

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 priorityqueue.
- 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 420

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

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

Answer: ACDF

Explanation:

http://media.amazonwebservices.com/AWS_Security_Best_Practices.pdf

NEW QUESTION 425

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

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

Answer: D

NEW QUESTION 430

What is one key difference between an Amazon EBS-backed and an instance-store backed instance?

- A. Amazon EBS-backed instances can be stopped and restarted.
- B. Instance-store backed instances can be stopped and restarted.
- C. Auto scaling requires using Amazon EBS-backed instances.
- D. Virtual Private Cloud requires EBS backed instance

Answer: A

NEW QUESTION 431

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

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

Answer: A

Explanation:

Ensure that instances in your subnet have public IP addresses or Elastic IP addresses.
https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Internet_Gateway.html

NEW QUESTION 433

You launch an Amazon EC2 instance without an assigned AWS identity and Access Management (IAM) role. Later, you decide that the instance should be running with an IAM role. Which action must you take in order to have a running Amazon EC2 instance with an IAM role assigned to it?

- A. Create an image of the instance, and register the image with an IAM role assigned and an Amazon EBS volume mapping.
- B. Create a new IAM role with the same permissions as an existing IAM role, and assign it to the running instance.
- C. Create an image of the instance, add a new IAM role with the same permissions as the desired IAM role, and deregister the image with the new role assigned.
- D. Create an image of the instance, and use this image to launch a new instance with the desired IAM role assigned.

Answer: D

NEW QUESTION 435

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 438

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 441

What is a placement group?

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

Answer: B

Explanation:

A placement group is a logical grouping of instances within a single Availability Zone. Using placement groups enables applications to participate in a low-latency, 10 Gigabits per second (Gbps) network. Placement groups are recommended for applications that benefit from low network latency, high network throughput, or both.

NEW QUESTION 446

A company has a workflow that sends video files from their on-premise system to AWS for transcoding. They use EC2 worker instances that pull transcoding jobs from SQS. Why is SQS an appropriate service for this scenario?

- A. SQS guarantees the order of the messages.
- B. SQS synchronously provides transcoding output.
- C. SQS checks the health of the worker instances.
- D. SQS helps to facilitate horizontal scaling of encoding task

Answer: D

Explanation:

Imho the idea for SQS is to improve scalability.
Elastic Beanstalk is checking the health of EC2 instances, not sure if SQS does.

D. SQS helps to facilitate horizontal scaling of encoding tasks.

Yes, this is a great scenario for SQS. "Horizontal scaling" means you have multiple instances involved in the workload (encoding tasks in this case). You can drop messages indicating an encoding job needs to be performed into an SQS queue, immediately making the job notification message accessible to any number of encoding worker instances.

NEW QUESTION 447

When creation of an EBS snapshot is initiated, but not completed, the EBS volume:

- A. Can be used while the snapshot is in progress.
- B. Cannot be detached or attached to an EC2 instance until the snapshot completes
- C. Can be used in read-only mode while the snapshot is in progress.
- D. Cannot be used until the snapshot complete

Answer: A

Explanation:

Snapshots occur asynchronously; the point-in-time snapshot is created immediately, but the status of the snapshot is pending until the snapshot is complete (when all of the modified blocks have been transferred to Amazon S3), which can take several hours for large initial snapshots or subsequent snapshots where many blocks have changed. While it is completing, an in-progress snapshot is not affected by ongoing reads and writes to the volume.
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-creating-snapshot.html>

NEW QUESTION 452

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

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

Answer: CE

NEW QUESTION 453

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 456

How can you secure data at rest on an EBS volume?

- A. Attach the volume to an instance using EC2's SSL interface.
- B. Write the data randomly instead of sequentially.
- C. Encrypt the volume using the S3 server-side encryption service.
- D. Create an IAM policy that restricts read and write access to the volume.
- E. Use an encrypted file system on top of the EBS volume

Answer: E

NEW QUESTION 460

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 462

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

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

Answer: B

Explanation:

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

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

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

NEW QUESTION 464

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

- A. Virtual Private Network connection
- B. AWS Directory Services, and ClassicLink
- C. Virtual Private Network connection
- D. AWS Directory Services, and Amazon Workspaces
- E. AWS Directory Service, Amazon Workspaces, and AWS Identity and Access Management
- F. Amazon Elastic Compute Cloud, and AWS Identity and Access Management

Answer: B

Explanation:

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

NEW QUESTION 468

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

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

Answer: CE

Explanation:

The screenshot shows a question: "Q: What data consistency model does Amazon S3 employ?". The answer text states: "Amazon S3 buckets in all Regions provide read-after-write consistency for PUTS of new objects and eventual consistency for overwrite PUTS and DELETES." Below this, there is a section titled "PUT Object" with a "Description" that says: "This implementation of the PUT operation adds an object to a bucket. You must have WRITE permissions on a bucket to add an object to it." Another line of text states: "Amazon S3 never adds partial objects. If you receive a success response, Amazon S3 added the entire object to the bucket."

NEW QUESTION 471

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

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

Answer: AC

Explanation:

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

NEW QUESTION 475

Which of the following are characteristics of a reserved instance? (Choose three.)

- A. It can be migrated across Availability Zones
- B. It is specific to an Amazon Machine Image (AMI)
- C. It can be applied to instances launched by Auto Scaling
- D. It is specific to an instance Type
- E. It can be used to lower Total Cost of Ownership (TCO) of a system

Answer: ACE

Explanation:

You can use Auto Scaling or other AWS services to launch the On-Demand instances that use your Reserved Instance benefits. For information about launching On-Demand instances, see Launch Your Instance. For information about launching instances using Auto Scaling, see the Auto Scaling User Guide.
<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/concepts-on-demand-reservedinstances.html>
<https://forums.aws.amazon.com/thread.jspa?threadID=56501>

NEW QUESTION 480

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

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

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