



Amazon-Web-Services

Exam Questions SAA-C01

AWS Certified Solutions Architect - Associate

NEW QUESTION 1

Your customer wishes to deploy an enterprise application to AWS which will consist of several web servers, several application servers and a small (50GB) Oracle database information is stored, both in the database and the file systems of the various servers. The backup system must support database recovery whole server and whole disk restores, and individual file restores with a recovery time of no more than two hours. They have chosen to use RDS Oracle as the database. Which backup architecture will meet these requirements?

- A. Backup RDS using automated daily DB backups Backup the EC2 instances using AMIs and supplement with file-level backup to S3 using traditional enterprise backup software to provide file level restore
- B. Backup RDS using a Multi-AZ Deployment Backup the EC2 instances using Amis, and supplement by copying file system data to S3 to provide file level restore.
- C. Backup RDS using automated daily DB backups Backup the EC2 instances using EBS snapshots and supplement with file-level backups to Amazon Glacier using traditional enterprise backup software to provide file level restore
- D. Backup RDS database to S3 using Oracle RMAN Backup the EC2 instances using Amis, and supplement with EBS snapshots for individual volume restore.

Answer: A

Explanation:

You need to use enterprise backup software to provide file level restore. See

https://d0.awsstatic.com/whitepapers/Backup_and_Recovery_Approaches_Using_AWS.pdf Page 18:

If your existing backup software does not natively support the AWS cloud, you can use AWS storage gateway products. AWS Storage Gateway is a virtual appliance that provides seamless and secure integration between your data center and the AWS storage infrastructure.

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 stale Eventual consistency) rather than an ACID (Atomicity. Consistency isolation. Durability) consistency model. The application is exhibiting undesirable behavior because the database is not able to handle the volume of writes. How can you reduce the load on your on-premises database resources in the most cost-effective way?

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

Answer: B

NEW QUESTION 3

Company B is launching a new game app for mobile devices. Users will log into the game using their existing social media account to streamline data capture. Company B would like to directly save player data and scoring information from the mobile app to a DynamoDS table named Score Data When a user saves their game the progress data will be stored to the Game state S3 bucket. What is the best approach for storing data to DynamoDB and S3?

- A. Use an EC2 Instance that is launched with an EC2 role providing access to the Score Data DynamoDB table and the GameState S3 bucket that communicates with the mobile app via web services.
- B. Use temporary security credentials that assume a role providing access to the Score Data DynamoDB table and the Game State S3 bucket using web identity federation.
- C. Use Login with Amazon allowing users to sign in with an Amazon account providing the mobile app with access to the Score Data DynamoDB table and the Game State S3 bucket.
- D. Use an IAM user with access credentials assigned a role providing access to the Score Data DynamoDB table and the Game State S3 bucket for distribution with the mobile app.

Answer: B

Explanation:

The requirements state “Users will log into the game using their existing social media account to streamline data capture.” This is what Cognito is used for, ie Web Identity Federation. Amazon also recommend to “build your app so that it requests temporary AWS security credentials dynamically when needed using web identity federation.”

NEW QUESTION 4

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 5

Your company is in the process of developing a next generation pet collar that collects biometric information to assist families with promoting healthy lifestyles for their pets Each collar will push 30kb of biometric data In JSON format every 2 seconds to a collection platform that will process and analyze the data providing health trending information back to the pet owners and veterinarians via a web portal Management has tasked you to architect the collection platform ensuring the following requirements are met. Provide the ability for real-time analytics of the inbound biometric data Ensure processing of the biometric data is highly durable. Elastic and parallel The results of

the analytic processing should be persisted for data mining
Which architecture outlined below win meet the initial requirements for the collection platform?

- A. Utilize S3 to collect the inbound sensor data analyze the data from S3 with a daily scheduled Data Pipeline and save the results to a Redshift Cluster.
- B. Utilize Amazon Kinesis to collect the inbound sensor data, analyze the data with Kinesis clients and save the results to a Redshift cluster using EMR.
- C. Utilize SQS to collect the inbound sensor data analyze the data from SQS with Amazon Kinesis and save the results to a Microsoft SQL Server RDS instance.
- D. Utilize EMR to collect the inbound sensor data, analyze the data from EUR with Amazon Kinesis and save me results to DynamoDB.

Answer: B

Explanation:

The POC solution is being scaled up by 1000, which means it will require 72TB of Storage to retain 24 months' worth of data. This rules out RDS as a possible DB solution which leaves you with RedShift. I believe DynamoDB is a more cost effective and scales better for ingest rather than using EC2 in an auto scaling group. Also, this example solution from AWS is somewhat similar for reference.
http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_timeseriesprocessing_16.pdf

NEW QUESTION 6

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 SOL 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 7

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

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

Answer: AB

Explanation:

The Route 53 and IAM are global.

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

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

NEW QUESTION 8

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

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

Answer: D

Explanation:

Amazon RDS Multi-AZ Deployments

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

Enhanced Durability

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

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

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

Increased Availability

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

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

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

On instance failure in Amazon Aurora deployments, Amazon RDS uses RDS Multi-AZ technology to automate failover to one of up to 15 Amazon Aurora Replicas you have created in any of three Availability Zones. If no Amazon Aurora Replicas have been provisioned, in the case of a failure, Amazon RDS will attempt to create a new Amazon Aurora DB instance for you automatically. <https://www.airpair.com/aws/posts/building-a-scalable-web-app-on-amazon-web-services-p1>

NEW QUESTION 9

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 10

Your company currently has a 2-tier web application running in an on-premises data center. You have experienced several infrastructure failures in the past two months resulting in significant financial losses. Your CIO is strongly agreeing to move the application to AWS. While working on achieving buy-in from the other company executives, he asks you to develop a disaster recovery plan to help improve Business continuity in the short term. He specifies a target Recovery Time Objective (RTO) of 4 hours and a Recovery Point Objective (RPO) of 1 hour or less. He also asks you to implement the solution within 2 weeks. Your database is 200GB in size and you have a 20Mbps Internet connection. How would you do this while minimizing costs?

- A. Create an EBS backed private AMI which includes a fresh install of your application
- B. Develop a CloudFormation template which includes your AMI and the required EC2, AutoScaling, and ELB resources to support deploying the application across Multiple- Availability-Zone
- C. Asynchronously replicate transactions from your on-premises database to a database instance in AWS across a secure VPN connection.
- D. Deploy your application on EC2 instances within an Auto Scaling group across multiple availability zone
- E. Asynchronously replicate transactions from your on-premises database to a database instance in AWS across a secure VPN connection.
- F. Create an EBS backed private AMI which includes a fresh install of your application
- G. Setup a script in your data center to backup the local database every 1 hour and to encrypt and copy the resulting file to an S3 bucket using multi-part upload.
- H. Install your application on a compute-optimized EC2 instance capable of supporting the application's average load
- I. Synchronously replicate transactions from your on-premises database to a database instance in AWS across a secure Direct Connect connection.

Answer: A

Explanation:

Overview of Creating Amazon EBS-Backed AMIs

First, launch an instance from an AMI that's similar to the AMI that you'd like to create. You can

connect to your instance and customize it. When the instance is configured correctly, ensure data integrity by stopping the instance before you create an AMI, then create the image. When you create an Amazon EBS-backed AMI, we automatically register it for you.

Amazon EC2 powers down the instance before creating the AMI to ensure that everything on the instance is stopped and in a consistent state during the creation process. If you're confident that your instance is in a consistent state appropriate for AMI creation, you can tell Amazon EC2 not to power down and reboot the

instance. Some file systems, such as XFS, can freeze and unfreeze activity, making it safe to create the image without rebooting the instance. During the AMI-creation process, Amazon EC2 creates snapshots of your instance's root volume and any other EBS volumes attached to your instance. If any volumes attached to the instance are encrypted, the new AMI only launches successfully on instances that support Amazon EBS encryption. For more information, see Amazon EBS Encryption.

Depending on the size of the volumes, it can take several minutes for the AMI-creation process to complete (sometimes up to 24 hours). You may find it more efficient to create snapshots of your volumes prior to creating your AMI. This way, only small, incremental snapshots need to be created when the AMI is created, and the process completes more quickly (the total time for snapshot creation remains the same). For more information, see Creating an Amazon EBS Snapshot. After the process completes, you have a new AMI and snapshot created from the root volume of the instance. When you launch an instance using the new AMI, we create a new EBS volume for its root volume using the snapshot. Both the AMI and the snapshot incur charges to your account until you delete them. For more information, see Deregistering Your AMI.

If you add instance-store volumes or EBS volumes to your instance in addition to the root device volume, the block device mapping for the new AMI contains information for these volumes, and the block device mappings for instances that you launch from the new AMI automatically contain information for these volumes. The instance-store volumes specified in the block device mapping for the new instance are new and don't contain any data from the instance store volumes of the instance you used to create the AMI. The data on EBS volumes persists. For more information, see Block Device Mapping.

NEW QUESTION 10

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

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

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 11

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

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

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

Answer: A

Explanation:

http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_mediasharing_09.pdf

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

NEW QUESTION 16

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

How should they architect their solution?

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

Answer: A

Explanation:

B is incorrect as you do not have insight into the public IP associated with a VPC Internet Gateway. C is incorrect as ELB receives a public DNS name.

D would exceed the maximum of 4 whitelisting IP addresses.

NEW QUESTION 17

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 lin
- 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 lin
- 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 lin
- 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 19

You are tasked with moving a legacy application from a virtual machine running Inside your datacenter to an Amazon VPC. Unfortunately this app requires access to a number of on-premises services and no one who configured the app still works for your company. Even worse there's no documentation for it. What will allow the application running inside the VPC to reach back and access its internal dependencies without being reconfigured? (Choose three.)

- A. An AWS Direct Connect link between the VPC and the network housing the internal services.
- B. An Internet Gateway to allow a VPN connection.
- C. An Elastic IP address on the VPC instance
- D. An IP address space that does not conflict with the one on-premises
- E. Entries in Amazon Route 53 that allow the Instance to resolve its dependencies' IP addresses
- F. A VM Import of the current virtual machine

Answer: ADF

Explanation:

AWS Direct Connect

AWS Direct Connect makes it easy to establish a dedicated network connection from your premises to AWS. Using AWS Direct Connect, you can establish private connectivity between AWS and your datacenter, office, or colocation environment, which in many cases can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than Internetbased connections.

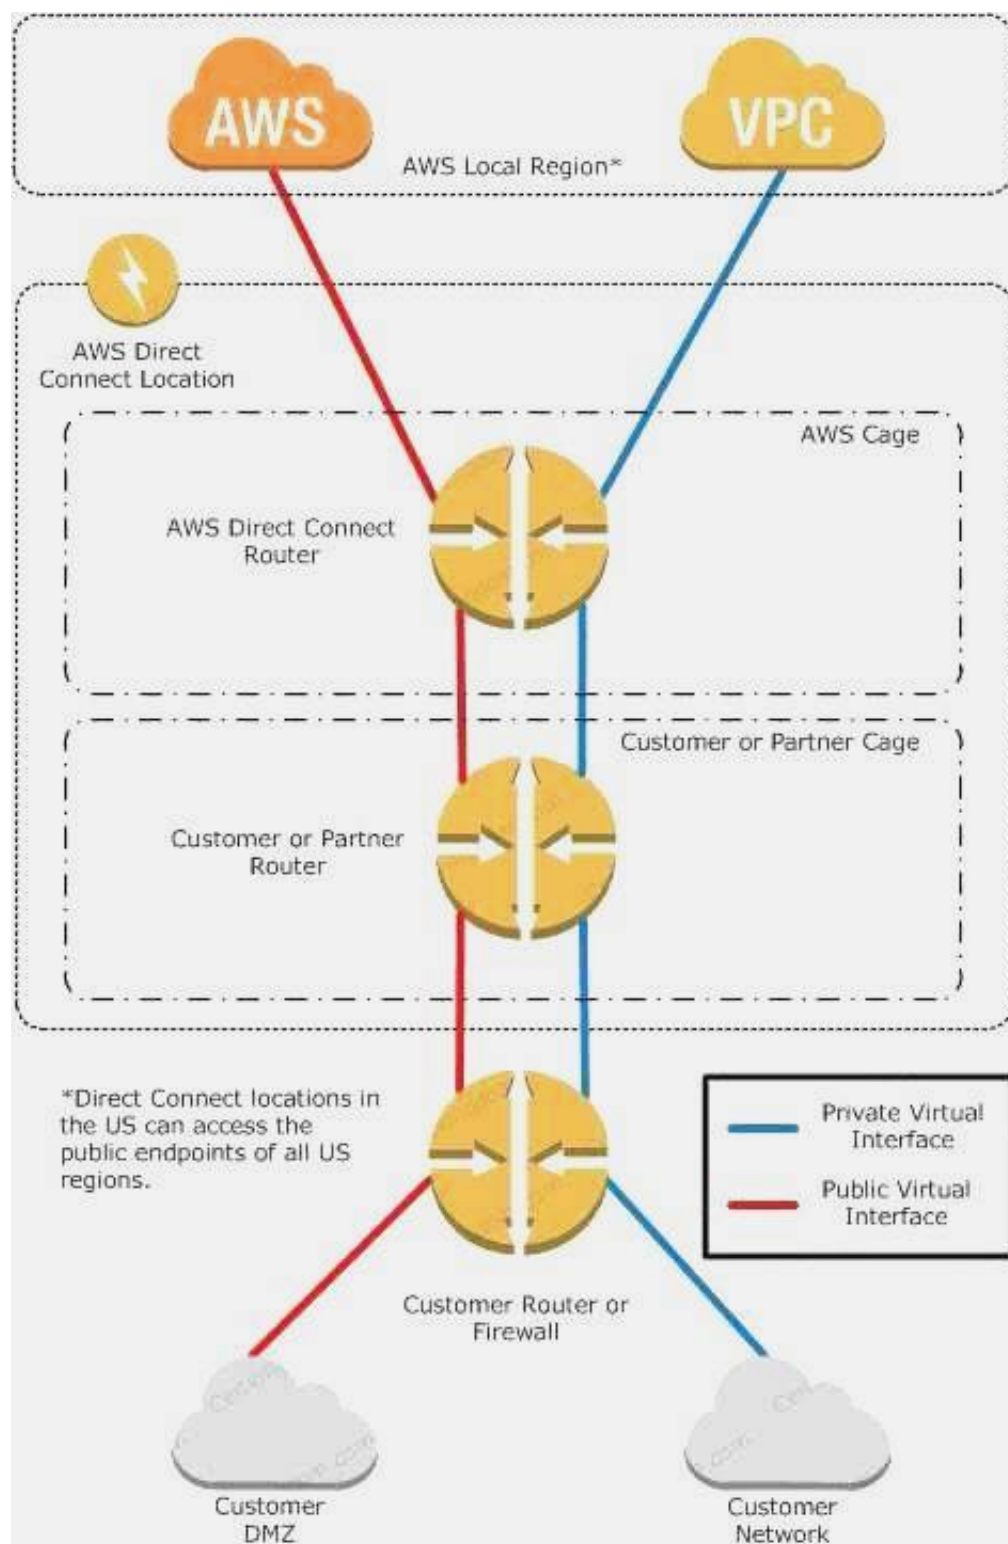
AWS Direct Connect lets you establish a dedicated network connection between your network and one of the AWS Direct Connect locations. Using industry standard 802.1q VLANs, this dedicated connection can be partitioned into multiple virtual interfaces. This allows you to use the same connection to access public resources such as objects stored in Amazon S3 using public IP address space, and private resources such as Amazon EC2 instances running within an Amazon Virtual Private Cloud (VPC) using private IP space, while maintaining network separation between the public and private environments. Virtual interfaces can be reconfigured at any time to meet your changing needs.

What is AWS Direct Connect?

AWS Direct Connect links your internal network to an AWS Direct Connect location over a standard 1 gigabit or 10 gigabit Ethernet fiber-optic cable. One end of the cable is connected to your router, the other to an AWS Direct Connect router. With this connection in place, you can create virtual interfaces directly to the AWS cloud (for example, to Amazon Elastic Compute Cloud (Amazon EC2) and Amazon Simple Storage Service (Amazon S3)) and to Amazon Virtual Private Cloud (Amazon VPC), bypassing Internet service providers in your network path. An AWS Direct Connect location provides access to Amazon Web Services in the region it is associated with, as well as access to other

US regions. For example, you can provision a single connection to any AWS Direct Connect location in the US and use it to access public AWS services in all US Regions and AWS GovCloud (US).

The following diagram shows how AWS Direct Connect interfaces with your network.



Requirements

To use AWS Direct Connect, your network must meet one of the following conditions:

Your network is colocated with an existing AWS Direct Connect location. For more information on available AWS Direct Connect locations, go to <http://aws.amazon.com/directconnect/>.

You are working with an AWS Direct Connect partner who is a member of the AWS Partner Network (APN). For a list of AWS Direct Connect partners who can help you connect, go to <http://aws.amazon.com/directconnect>.

You are working with an independent service provider to connect to AWS Direct Connect. In addition, your network must meet the following conditions:

Connections to AWS Direct Connect require single mode fiber, 1000BASE-LX (1310nm) for 1 gigabit Ethernet, or 10GBASE-LR (1310nm) for 10 gigabit Ethernet. Auto Negotiation for the port must be disabled. You must support 802.1Q VLANs across these connections.

Your network must support Border Gateway Protocol (BGP) and BGP MD5 authentication. Optionally, you may configure Bidirectional Forwarding Detection (BFD).

To connect to Amazon Virtual Private Cloud (Amazon VPC), you must first do the following: Provide a private Autonomous System Number (ASN). Amazon allocates a private IP address in the 169.x.x.x range to you.

Create a virtual private gateway and attach it to your VPC. For more information about creating a virtual private gateway, see Adding a Hardware Virtual Private Gateway to Your VPC in the Amazon VPC User Guide.

To connect to public AWS products such as Amazon EC2 and Amazon S3, you need to provide the following:

A public ASN that you own (preferred) or a private ASN.

Public IP addresses (/31) (that is, one for each end of the BGP session) for each BGP session. If you do not have public IP addresses to assign to this connection, log on to AWS and then open a ticket with AWS Support.

The public routes that you will advertise over BGP.

NEW QUESTION 21

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

A. Use S3 with reduced redundancy to store and serve the scanned files, install the commercial search application on EC2 Instances and configure with auto-scaling and an Elastic Load Balancer.

B. Model the environment using CloudFormation use an EC2 instance running Apache webserver and an open source search application, stripe multiple standard EBS volumes together to store the JPEGs and search index.

C. Use S3 with standard redundancy to store and serve the scanned files, use CloudSearch for query processing, and use Elastic Beanstalk to host the website across multiple availability zones.

D. Use a single-AZ RDS MySQL instance to store the search index and the JPEG images use an EC2 instance to serve the website and translate user queries into SQL.

E. Use a CloudFront download distribution to serve the JPEGs to the end users and Install the current commercial search product, along with a Java Container Tomcat.

the website on EC2 instances and use Route53 with DNS round-robin.

Answer: C

Explanation:

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

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

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

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

EBS will be way more expensive, and you'll need 2x the drives if you need 2 boxes. Yuck. Consider a smaller storage format. For example, JPEG200 or WebP or other tools might make for smaller images. There is also the DejaVu format from a while back.

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

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

load if you do it right. Database Layer: A few options:

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

Use RDS to run MySQL/Postgres

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

App Layer:

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

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

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

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

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

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

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

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

We're missing some information like load, latency expectations from search, indexing speed, size of the search index, etc. But with what you've given us, I would go with S3 as the storage for the files (S3 rocks. It is really, really awesome). If you're stuck with the commercial search application, then on EC2 instances with autoscaling and an ELB. If you are allowed an alternative search engine, Elasticsearch is probably your best bet. I'd run it on EC2 instead of the AWS

Elasticsearch service, as IMHO it's not ready yet. Don't autoscale Elasticsearch automatically though, it'll cause all sorts of issues. I have zero experience with CloudSearch so I can't comment on that. Regardless of which option, I'd use CloudFormation for all of it.

NEW QUESTION 22

A corporate web application is deployed within an Amazon Virtual Private Cloud (VPC) and is connected to the corporate data center via an IPsec VPN. The application must authenticate against the on-premises LDAP server. After authentication, each logged-in user can only access an Amazon Simple Storage Space (S3) keyspace specific to that user.

Which two approaches can satisfy these objectives? (Choose two.)

- A. Develop an identity broker that authenticates against IAM security Token service to assume a IAM role in order to get temporary AWS security credentials The application calls the identity broker to get AWS temporary security credentials with access to the appropriate S3 bucket.
- B. The application authenticates against LDAP and retrieves the name of an IAM role associated with the user
- C. The application then calls the IAM Security Token Service to assume that IAM role
- D. The application can use the temporary credentials to access the appropriate S3 bucket.
- E. Develop an identity broker that authenticates against LDAP and then calls IAM Security Token Service to get IAM federated user credential
- F. The application calls the identity broker to get IAM federated user credentials with access to the appropriate S3 bucket.
- G. The application authenticates against LDAP the application, then calls the AWS identity and Access Management (IAM) Security service to log in to IAM using the LDAP credentials, the application can use the IAM temporary credentials to access the appropriate S3 bucket.
- H. The application authenticates against IAM Security Token Service using the LDAP credentials, the application uses those temporary AWS security credentials to access the appropriate S3 bucket.

Answer: BC

Explanation:

Imagine that in your organization, you want to provide a way for users to copy data from their computers to a backup folder. You build an application that users can run on their computers. On the back end, the application reads and writes objects in an S3 bucket. Users don't have direct access to AWS. Instead, the application communicates with an identity provider (IdP) to authenticate the user. The IdP gets the user information from your organization's identity store (such as an LDAP directory) and then generates a SAML assertion that includes authentication and authorization information about that user. The application then uses that assertion to make a call to the AssumeRoleWithSAML API to get temporary security credentials. The app can then use those credentials to access a folder in the S3 bucket that's specific to the user. http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_saml.html

NEW QUESTION 27

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 block-based 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 29

You are the new IT architect in a company that operates a mobile sleep tracking application When activated at night, the mobile app is sending collected data points of 1 kilobyte every 5 minutes to your backend

The backend takes care of authenticating the user and writing the data points into an Amazon DynamoDB table.

Every morning, you scan the table to extract and aggregate last night's data on a per user basis, and store the results in Amazon S3.

Users are notified via Amazon SMS mobile push notifications that new data is available, which is parsed and visualized by (The mobile app Currently you have around 100k users who are mostly based out of North America.

You have been tasked to optimize the architecture of the backend system to lower cost what would you recommend? (Choose two.)

- A. Create a new Amazon DynamoDB (able each day and drop the one for the previous day after its data is on Amazon S3.
- B. Have the mobile app access Amazon DynamoDB directly instead of JSON files stored on Amazon S3.
- C. Introduce an Amazon SQS queue to buffer writes to the Amazon DynamoDB table and reduce provisioned write throughput.
- D. Introduce Amazon ElastiCache to cache reads from the Amazon DynamoDB table and reduce provisioned read throughput.
- E. Write data directly into an Amazon Redshift cluster replacing both Amazon DynamoDB and Amazon S3.

Answer: AC

NEW QUESTION 30

You've been hired to enhance the overall security posture for a very large e-commerce site. They

have a well architected multi-tier application running in a VPC that uses ELBs in front of both the web and the app tier with static assets served directly from S3

They are using a combination of RDS and DynamoDB for their dynamic data and then archiving nightly into S3 for further processing with EMR They are

concerned because they found questionable log entries and suspect someone is attempting

to gain unauthorized access.

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

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

Answer: C

NEW QUESTION 33

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 36

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

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

Answer: C

NEW QUESTION 38

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

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

Answer: C

NEW QUESTION 43

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

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

Answer: AD

Explanation:

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

NEW QUESTION 48

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 49

You are designing a social media site and are considering how to mitigate distributed denial-of-service (DDoS) attacks. Which of the below are viable mitigation techniques? (Choose three.)

- A. Add multiple elastic network interfaces (ENIs) to each EC2 instance to increase the network bandwidth.
- B. Use dedicated instances to ensure that each instance has the maximum performance possible.
- C. Use an Amazon CloudFront distribution for both static and dynamic content.
- D. Use an Elastic Load Balancer with auto scaling groups at the web tier.
- E. App and Amazon Relational Database Service (RDS) tiers.
- F. Add alert Amazon CloudWatch to look for high Network in and CPU utilization.
- G. Create processes and capabilities to quickly add and remove rules to the instance OS firewall.

Answer: CDE

NEW QUESTION 54

A benefits enrollment company is hosting a 3-tier web application running in a VPC on AWS which includes a NAT (Network Address Translation) instance in the public Web tier. There is enough provisioned capacity for the expected workload for the new fiscal year benefit enrollment period plus some extra overhead. Enrollment proceeds nicely for two days and then the web tier becomes unresponsive. Upon investigation using CloudWatch and other monitoring tools, it is discovered that there is an extremely large and unanticipated amount of inbound traffic coming from a set of 15 specific IP addresses over port 80 from a country where the benefits company has no customers. The web tier instances are so overloaded that benefit enrollment administrators cannot even SSH into them. Which activity would be useful in defending against this attack?

- A. Create a custom route table associated with the web tier and block the attacking IP addresses from the IGW (Internet Gateway).
- B. Change the EIP (Elastic IP Address) of the NAT instance in the web tier subnet and update the Main Route Table with the new EIP.
- C. Create 15 Security Group rules to block the attacking IP addresses over port 80.
- D. Create an inbound NACL (Network Access Control List) associated with the web tier subnet with deny rules to block the attacking IP addresses.

Answer: D

Explanation:

Use AWS Identity and Access Management (IAM) to control who in your organization has permission to create and manage security groups and network ACLs (NACL). Isolate the responsibilities and roles for better defense. For example, you can give only your network administrators or security admin the permission to manage the security groups and restrict other roles.

NEW QUESTION 56

Your fortune 500 company has under taken a TCO analysis evaluating the use of Amazon S3 versus acquiring more hardware. The outcome was that all employees would be granted access to use Amazon S3 for storage of their personal documents. Which of the following will you need to consider so you can set up a solution that incorporates single sign-on from your corporate AD or LDAP directory and restricts access for each user to a designated user folder in a bucket? (Choose three.)

- A. Setting up a federation proxy or identity provider
- B. Using AWS Security Token Service to generate temporary tokens
- C. Tagging each folder in the bucket
- D. Configuring IAM role
- E. Setting up a matching IAM user for every user in your corporate directory that needs access to a folder in the bucket

Answer: ABD

NEW QUESTION 57

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

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

Answer: ACD

Explanation:

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

NEW QUESTION 61

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 CloudFront distribution with Restrict Viewer Access Forward Query string set to true and minimum TTL of 0.

Answer: C

NEW QUESTION 66

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

- A. Deploy ElasticCache in-memory cache running in each availability zone
- B. Implement sharding to distribute load to multiple RDS MySQL instances
- C. Increase the RDS MySQL Instance size and Implement provisioned IOPS
- D. Add an RDS MySQL read replica in each availability zone

Answer: AD

NEW QUESTION 67

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 CloudWatch
- B. And RDS with 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 CloudWatch
- 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 70

You are developing a new mobile application and are considering storing user preferences in AWS. This would provide a more uniform cross-device experience to users using multiple mobile devices to access the application. The preference data for each user is estimated to be 50KB in size. Additionally, 5 million customers are expected to use the application on a regular basis. The solution needs to be cost-effective, highly available, scalable and secure. How would you design a solution to meet the above requirements?

- A. Setup an RDS MySQL instance in 2 availability zones to store the user preference data
- B. Deploy a public facing application on a server in front of the database to manage security and access credentials
- C. Setup a DynamoDB table with an item for each user having the necessary attributes to hold the user preference
- D. The mobile application will query the user preferences directly from the DynamoDB table
- E. Utilize STS
- F. Web Identity Federation, and DynamoDB Fine Grained Access Control to authenticate and authorize access.
- G. Setup an RDS MySQL instance with multiple read replicas in 2 availability zones to store the user preference data. The mobile application will query the user preferences from the read replica
- H. Leverage the MySQL user management and access privilege system to manage security and access credentials.
- I. Store the user preference data in S3. Setup a DynamoDB table with an item for each user and an item attribute pointing to the user's S3 object
- J. The mobile application will retrieve the S3 URL from DynamoDB and then access the S3 object directly. Utilize STS, Web identity Federation, and S3 ACLs to authenticate and authorize access.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/aws/fine-grained-access-control-for-amazon-dynamodb/> Here are some of the things that you can build using fine-grained access control:

A mobile app that displays information for nearby airports, based on the user's location. The app can access and display attributes such as airline names, arrival times, and flight numbers. However, it cannot access or display pilot names or passenger counts.

A mobile game which stores high scores for all users in a single table. Each user can update their own scores, but has no access to the other ones.

NEW QUESTION 74

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 variable
- 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 app'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 77

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 account's bill to a Master AWS account using Consolidated Billing. To make sure you keep within budget, you would like to implement a way for administrators in the Master account to have access to stop, delete and/or terminate resources in both the Dev and Test accounts. Identify which option will allow you to achieve this goal.

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

Answer: C

NEW QUESTION 80

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

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

- A. Send all the log events to Amazon SQS

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

Answer: B

Explanation:

Amazon Kinesis Streams allows for real-time data processing. With Amazon Kinesis Streams, you can continuously collect data as it is generated and promptly react to critical information about your business and operations.
<https://aws.amazon.com/kinesis/streams/>

NEW QUESTION 83

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

- A. Configure an instance with monitoring software and the elastic network interface (ENI) set to promiscuous mode packet sniffing to see an traffic across the VPC,
- B. Create a second VPC and route all traffic from the primary application VPC through the second VPC where the scalable virtualized IDS/IPS platform resides.
- C. Configure servers running in the VPC using the host-based 'route' commands to send all traffic through the platform to a scalable virtualized IDS/IPS.
- D. Configure each host with an agent that collects all network traffic and sends that traffic to the IDS/IPS platform for inspection.

Answer: D

Explanation:

\A. Not possible to set an instance's NIC into promiscuous mode.

\B. Incorrect... VPC peering connections are not "transitive", i.e. you cannot pass traffic through a VPC peering connection into another VPC, and then have that other VPC send the traffic to some third VPC, or the Internet, or a VPN, or a direct connect circuit. (I would assume AWS does not allow redistribution of routes from one VPC's back-end VRF into another VPC's back-end VRF, unless it is that first VPC's CIDR block? Someone from AWS would have to chime in here, and they're probably not going to tell us.)

\C. This one is incorrect because adding static routes on an instance won't affect the routing from any point after the packet leaves the instance's NIC. AWS will check the destination IP address in the packet header and forward according to the VPC routing table's routes. You'd need to make routing changes in the VPC route table for that instance's traffic to get sent through another device (e.g. NAT gateway, VPN instance, or security proxy in this case). (You could tunnel/proxy the traffic over through the IPS tier by changing the destination IP address in the IP header of the packet before it left the instance. But choice C did not state anything about doing anything like that. It just said add a static route on the instance, which does not change the destination IP address in the IP header of the packet.)

\D. Correct, this is the standard approach, and is definitely scalable.

NEW QUESTION 86

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 87

Your firm has uploaded a large amount of aerial image data to S3. In the past, in your on-premises environment, you used a dedicated group of servers to oaten process this data and used Rabbit MQ - An open source messaging system to get job information to the servers. Once processed the data would go to tape and be shipped offsite. Your manager told you to stay with the current design, and leverage AWS archival storage and messaging services to minimize cost. Which is correct?

- A. Use SQS for passing job messages use Cloud Watch alarms to terminate EC2 worker instances when they become idl
- B. Once data is processed, change the storage class of the S3 objects to Reduced Redundancy Storage.
- C. Setup Auto-Scaled workers triggered by queue depth that use spot instances to process messages in SOS Once data is processed,
- D. Change the storage class of the S3 objects to Reduced Redundancy Storag
- E. Setup Auto-Scaled workers triggered by queue depth that use spot instances to process messages in SQS Once data is processed, change the storage class of the S3 objects to Glacier.
- F. Use SNS to pass job messages use Cloud Watch alarms to terminate spot worker instances when they become idl
- G. Once data is processed, change the storage class of the S3 object to Glacier.

Answer: C

Explanation:

The question key part to focus on is “and leverage AWS archival storage and messaging services to minimize cost.”

For that the storage that is the lowest cost in the answers is Glacier, in addition, the messaging cost is less for SQS then for SNS if they both exceed 1 million transactions which is free. The

only answer that satisfies the above two criteria is answer C. Also, there does not seem to be an urgency in speed of messaging therefore SQS satisfies that need. SNS being more real time delivery mechanism.

NEW QUESTION 92

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 95

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 99

What does RRS stand for when talking about S3?

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

Answer: D

Explanation:

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

NEW QUESTION 103

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 105

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

- A. Security Pool
- B. Secure Zone
- C. Security Token Pool
- D. Security Group

Answer: D

NEW QUESTION 110

When you run a DB Instance as a Multi-AZ deployment, the " " serves database writes and reads

- A. secondary
- B. backup
- C. stand by
- D. primary

Answer: D

NEW QUESTION 111

Every user you create in the IAM system starts with .

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

Answer: C

NEW QUESTION 116

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 119

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

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

Answer: C

NEW QUESTION 123

By default, EBS volumes that are created and attached to an instance at launch are deleted when that instance is terminated. You can modify this behavior by changing the value of the flag to false when you launch the instance

- A. DeleteOnTermination
- B. RemoveOnDeletion
- C. RemoveOnTermination
- D. TerminateOnDeletion

Answer: A

Explanation:

By default, Amazon EBS root device volumes are automatically deleted when the instance terminates. However, by default, any additional EBS volumes that you attach at launch, or any EBS volumes that you attach to an existing instance persist even after the instance terminates.

This behavior is controlled by the volume's DeleteOnTermination attribute, which you can modify.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/terminating-instances.html>

NEW QUESTION 125

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 130

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 135

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 140

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

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

Answer: A

NEW QUESTION 145

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

- A. TRUE
- B. FALSE

Answer: A

NEW QUESTION 147

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 150

What are the two permission types used by AWS?

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

Answer: D

NEW QUESTION 154

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 159

What is the maximum key length of a tag?

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

Answer: D

NEW QUESTION 163

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

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

Answer: D

NEW QUESTION 168

Which is the default region in AWS?

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

Answer: B

Explanation:

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

NEW QUESTION 173

What are the Amazon EC2 API tools?

- A. They don't exist
- B. The Amazon EC2 API 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 178

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

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

Answer: C

Explanation:

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

NEW QUESTION 180

Is creating a Read Replica of another Read Replica supported?

- A. Only in certain regions
- B. Only with MSSQL based RDS
- C. Only for Oracle RDS types
- D. No

Answer: D

Explanation:

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

Q: Can I create a Read Replica of another Read Replica?

Amazon RDS for MySQL and MariaDB: You can **create a second-tier Read Replica from an existing first-tier Read Replica**. By creating a second-tier Read Replica, you may be able to move some of the replication load from the master database instance to a first-tier Read Replica. Please note that a second-tier Read Replica may lag further behind the master because of additional replication latency introduced as transactions are replicated from the master to the first tier replica and then to the second-tier replica.

Amazon RDS for PostgreSQL: Read Replicas of Read Replicas are not currently supported.

NEW QUESTION 183

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 187

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 188

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 190

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

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

Answer: B

NEW QUESTION 195

How many types of block devices does Amazon EC2 support?

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

Answer: A

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/block-device-mapping-concepts.html> Amazon EC2 supports two types of block devices:
Instance store volumes (virtual devices whose underlying hardware is physically attached to the host computer for the instance)
EBS volumes (remote storage devices)

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

Block Device Mapping Concepts

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

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

NEW QUESTION 198

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 202

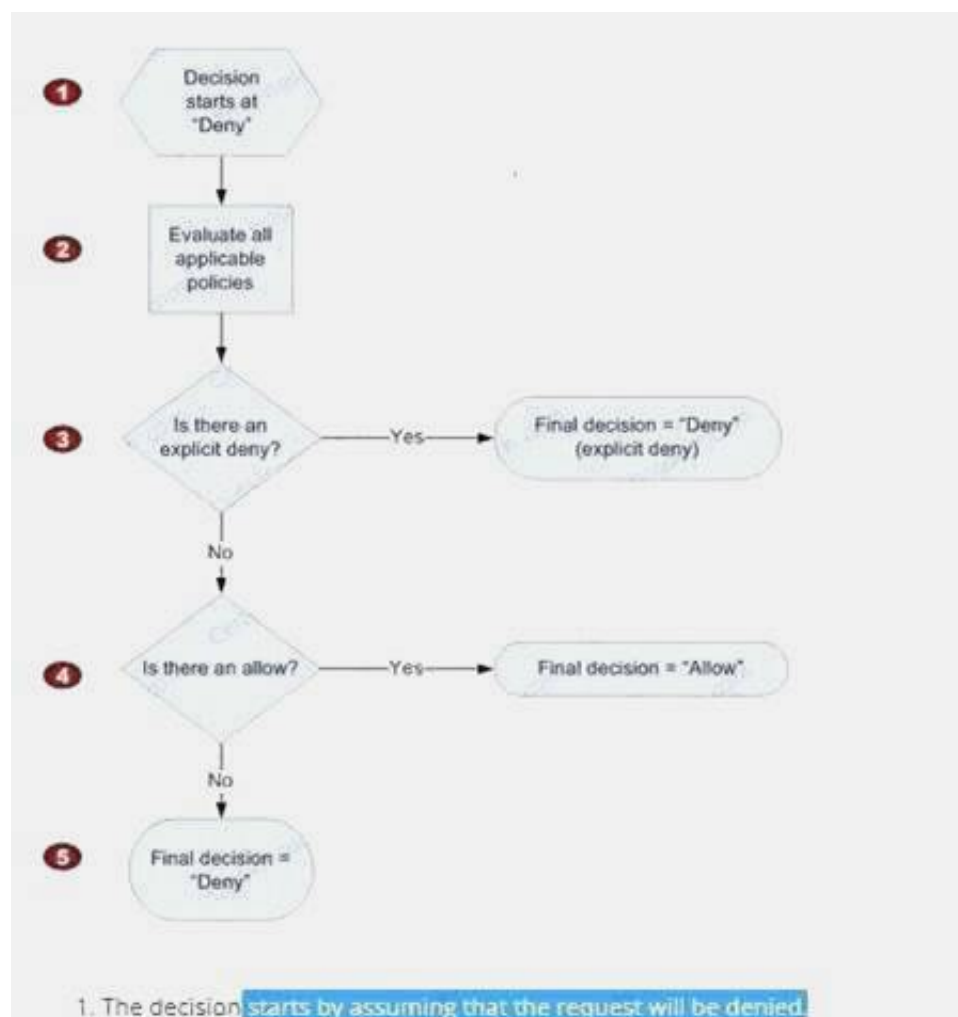
IAM's Policy Evaluation Logic always starts with a default ____ for every request, except for those that use the AWS account's root security credentials b

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

Answer: B

Explanation:

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



NEW QUESTION 207

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 210

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

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

Answer: B

Explanation:

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

NEW QUESTION 211

SQL Server _____ store logins and passwords in the master database.

- A. can be configured to but by default does not
- B. doesn't
- C. does

Answer: C

Explanation:

There are two authentications Windows authentication

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

SQL Server authentication.

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

NEW QUESTION 214

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

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

Answer: D

Explanation:

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



NEW QUESTION 215

What is an isolated database environment running in the cloud (Amazon RDS) called?

- A. DB Instance
- B. DB Server
- C. DB Unit
- D. DB Volume

Answer: A

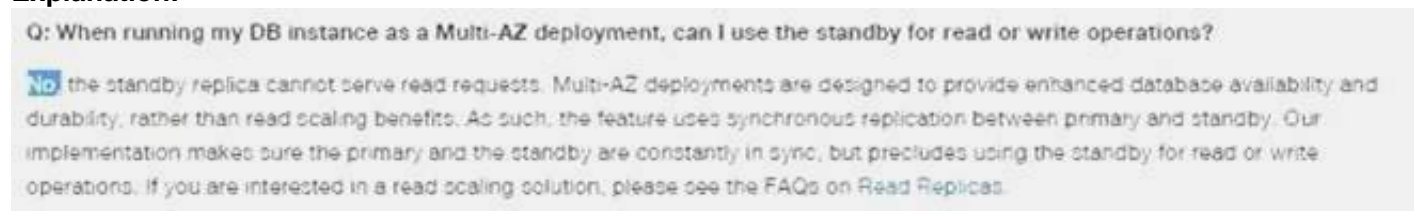
NEW QUESTION 216

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

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

Answer: D

Explanation:



NEW QUESTION 218

Which Amazon service can I use to define a virtual network that closely resembles a traditional data center?

- A. Amazon VPC
- B. Amazon ServiceBus
- C. Amazon EMR
- D. Amazon RDS

Answer: A

NEW QUESTION 221

What is the command line instruction for running the remote desktop client in Windows?

- A. desk.cpl
- B. mstsc

Answer: B

NEW QUESTION 222

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

- A. MyISAM
- B. InnoDB

Answer: B

NEW QUESTION 224

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 231

Making your snapshot public shares all snapshot data with everyone. Can the snapshots with AWS Marketplace product codes be made public?

- A. No
- B. Yes

Answer: A

Explanation:

"Making your snapshot public shares all snapshot data with everyone; however, snapshots with AWS Marketplace product codes cannot be made public.

Encrypted snapshots cannot be shared between

accounts or made public." <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-modifyingsnapshot-permissions.html>

"This is not a valid option for encrypted snapshots or snapshots with AWS Marketplace product codes."

NEW QUESTION 234

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 235

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 240

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 244

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 245

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 249

Can I delete a snapshot of the root device of an EBS volume used by a registered AMI?

- A. Only via API
- B. Only via Console
- C. Yes
- D. No

Answer: D

Explanation:

Note that you can't delete a snapshot of the root device of an EBS volume used by a registered AMI. You must first deregister the AMI before you can delete the snapshot.

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

NEW QUESTION 253

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 255

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

- A. FALSE
- B. TRUE

Answer: B

Explanation:

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

NEW QUESTION 259

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 260

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 261

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

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

Answer: B

Explanation:

Data persistence

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

References:

NEW QUESTION 262

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

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

Answer: A

Explanation:

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

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

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

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

Default Security Groups

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

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

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

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

You can't delete a default security group. If you try to delete the EC2-Classical default security group, you'll get the following error: `Client.InvalidGroup.Reserved: The security group 'default' is reserved.` If you try to delete a VPC default security group, you'll get the following error: `Client.CannotDelete: the specified group: "sg-51630134" name: "default" cannot be deleted by a user.`

NEW QUESTION 267

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 272

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 275

Please select the Amazon EC2 resource which can be tagged.

- A. key pairs
- B. Elastic IP addresses
- C. placement groups
- D. Amazon EBS snapshots

Answer: C

Explanation:

Placement group and Elastic IP cannot be tagged. http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html Snapshots can be tagged:

http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html

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
Reserved Instance listing	No	
Route table	Yes	None
Spot instance request	Yes	None
Security group - EC2-Classical	Yes	None
Security group - VPC	Yes	None
Snapshot	Yes	None

NEW QUESTION 279

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 280

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 281

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

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

Answer: B

Explanation:

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



NEW QUESTION 283

Is it possible to access your EBS snapshots?

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

Answer: B

Explanation:

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

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

NEW QUESTION 284

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

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

Answer: D

Explanation:

"Amazon EBS is particularly helpful for database-style applications that frequently encounter many random reads and writes across the data set."

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

NEW QUESTION 288

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

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

Answer: B

Explanation:

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

NEW QUESTION 290

Is decreasing the storage size of a DB Instance permitted?

- A. Depends on the RDMS used
- B. Yes
- C. No

Answer: C

Explanation:

"note that you cannot reduce storage size once it has been allocated" Source:

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuideCHAP_Storage.html#CHAP_Storage.FactsAbout

NEW QUESTION 293

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

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

Answer: A

NEW QUESTION 294

By default, what are ENIs that are automatically created and attached to instances using the EC2 console set to do when the attached instance terminates?

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

Answer: B

Explanation:

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

NEW QUESTION 297

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 301

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 304

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/>

Amazon EC2 Dedicated Hosts

An Amazon EC2 Dedicated Host is a physical server with EC2 instance capacity fully dedicated to your use. **Dedicated Hosts** can help you address compliance requirements and reduce costs by allowing you to use your existing server-bound software licenses. Visit the [EC2 Dedicated Host Pricing page](#) for information on availability and pricing.

NEW QUESTION 306

What does Amazon RDS stand for?

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

Answer: B

NEW QUESTION 307

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

- A. 30 minutes
- B. You always get instant responses (within a few seconds).
- C. 10 minutes
- D. 1 hour

Answer: D

NEW QUESTION 310

What does Amazon ELB stand for?

- A. Elastic Linux Box.
- B. Encrypted Linux Box.
- C. Encrypted Load Balancing.
- D. Elastic Load Balancin

Answer: D

NEW QUESTION 315

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 317

Does Dynamic DB support in-place atomic updates?

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

Answer: C

Explanation:

Q: Does DynamoDB support in-place atomic updates?

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

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

NEW QUESTION 322

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 326

What's an ECU?

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

Answer: B

Explanation:

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

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

NEW QUESTION 327

What does Amazon EBS stand for?

- A. Elastic Block Storage
- B. Elastic Business Server
- C. Elastic Blade Server
- D. Elastic Block Store

Answer: D

Explanation:

<https://aws.amazon.com/ebs/> Amazon Elastic Block Store (EBS)

Amazon Elastic Block Store (Amazon EBS) provides persistent block level storage volumes for use with Amazon EC2 instances in the AWS Cloud. Each Amazon EBS volume is automatically replicated within its Availability Zone to protect you from component failure, offering high availability and durability. Amazon EBS volumes offer the consistent and low- latency performance needed to run your workloads. With Amazon EBS, you can scale your usage up or down within minutes all while paying a low price for only what you provision.

Amazon Elastic Block **Store** (Amazon EBS) provides persistent block storage volumes for use with Amazon EC2 instances in the AWS Cloud. Each Amazon EBS volume is automatically replicated within its Availability Zone to protect you from component failure, offering high availability and durability. Amazon EBS volumes offer the consistent and low-latency performance needed to run your workloads. With Amazon EBS, you can scale your usage up or down within minutes – all while paying a low price for only what you provision.

NEW QUESTION 332

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 336

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 339

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 344

The Amazon EC2 web service can be accessed using the _____ web services messaging protocol. This interface is described by a Web Services Description Language (WSDL) document.

- A. SOAP
- B. DCOM
- C. CORBA
- D. XML-RPC

Answer: A

Explanation:

<http://docs.aws.amazon.com/AWSECommerceService/latest/DG/WSDLLocation.html>

WSDL and Schema Definitions

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

Note

Amazon RDS supports SOAP only through HTTPS.

NEW QUESTION 349

Is creating a Read Replica of another Read Replica supported?

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

Answer: D

NEW QUESTION 350

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 352

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 355

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

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

Answer: A

Explanation:

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

NEW QUESTION 357

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

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

Answer: A

Explanation:

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

NEW QUESTION 360

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

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

Answer: B

Explanation:

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

NEW QUESTION 361

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 366

Will I be alerted when automatic failover occurs?

- A. Only if SNS configured
- B. No
- C. Yes
- D. Only if Cloudwatch configured

Answer: A

Explanation:

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

Amazon RDS uses the Amazon Simple Notification Service (Amazon SNS) to provide notification when an Amazon RDS event occurs. These notifications can be in any notification form supported by Amazon SNS for an AWS region, such as an email, a text message, or a call to an HTTP endpoint. Amazon RDS groups these events into categories that you can subscribe to so that you can be notified when an event in that category occurs. C is not correct because even though event is created by RDS you will not be alerted for it unless you configure your subscription in SNS.

NEW QUESTION 369

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

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

Answer: BDF

Explanation:

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

NEW QUESTION 372

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 376

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

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

Answer: BE

NEW QUESTION 380

A company is building a two-tier web application to serve dynamic transaction-based content. The data tier is leveraging an Online Transactional Processing (OLTP) database. What services should you leverage to enable an elastic and scalable web tier?

- A. Elastic Load Balancing, Amazon EC2, and Auto Scaling
- B. Elastic Load Balancing, Amazon RDS with Multi-AZ, and Amazon S3
- C. Amazon RDS with Multi-AZ and Auto Scaling
- D. Amazon EC2, Amazon DynamoDB, and Amazon S3

Answer: A

NEW QUESTION 383

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 priorit
- E. Transformation instances first poll the high priority queue; if there is no message, they poll the default priorityqueue.
- F. Use a single SQS queu
- G. Each message contains the priority leve
- H. Transformation instances poll high-priority messages first.

Answer: C

NEW QUESTION 385

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

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

Answer: BD

Explanation:

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

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

NEW QUESTION 386

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 388

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

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 392

When you put objects in Amazon S3, what is the indication that an object was successfully stored?

- A. A HTTP 200 result code and MD5 checksum, taken together, indicate that the operation was successful.
- B. Amazon S3 is engineered for 99.999999999% durabilit
- C. Therefore there is no need to confirm that data was inserted.
- D. A success code is inserted into the S3 object metadata.
- E. Each S3 account has a special bucket named _s3_log
- F. Success codes are written to this bucket with a timestamp and checksum.

Answer: A

Explanation:

To ensure that data is not corrupted traversing the network, use the Content-MD5 form field. When you use this form field, Amazon S3 checks the object against the provided MD5 value. If they do not match, Amazon S3 returns an error. The status code returned to the client upon successful upload if

success_action_redirect is not specified. Accepts the values 200, 201, or 204 (default). <http://docs.aws.amazon.com/AmazonS3/latest/API/RESTObjectPOST.html>

NEW QUESTION 395

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 400

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

SAA-C01 Practice Exam Features:

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

100% Actual & Verified — Instant Download, Please Click
[Order The SAA-C01 Practice Test Here](#)