

Question: 1

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 3)

- A. Implement third party volume encryption tools
- B. Implement SSL/TLS for all services running on the server
- 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. Do nothing as EBS volumes are encrypted by default

Answer: A,C,D

Question: 2

A customer is deploying an SSL enabled web application to AWS and would like to implement a separation of roles between the EC2 service administrators that are entitled to login to instances as well as making API calls and the security officers who will maintain and have exclusive access to the application's X.509 certificate that contains the private key.

- A. Upload the certificate on an S3 bucket owned by the security officers and accessible only by EC2 Role of the web servers.
- B. Configure the web servers to retrieve the certificate upon boot from an CloudHSM is managed by the security officers.
- C. Configure system permissions on the web servers to restrict access to the certificate only to the authority security officers
- D. Configure IAM policies authorizing access to the certificate store only to the security officers and terminate SSL on an ELB.

Answer: D

Explanation:

You'll terminate the SSL at ELB. and the web request will get unencrypted to the EC2 instance, even if the certs are stored in S3, it has to be configured on the web servers or load balancers somehow, which becomes difficult if the keys are stored in S3. However, keeping the keys in the cert store and using IAM to restrict access gives a clear separation of concern between security officers and developers. Developer's personnel can still configure SSL on ELB without actually handling the keys.

Question: 3

You have recently joined a startup company building sensors to measure street noise and air quality

in urban areas. The company has been running a pilot deployment of around 100 sensors for 3 months each sensor uploads 1KB of sensor data every minute to a backend hosted on AWS. During the pilot, you measured a peak of 10 IOPS on the database, and you stored an average of 3GB of sensor data per month in the database.

The current deployment consists of a load-balanced auto scaled Ingestion layer using EC2 instances and a PostgreSQL RDS database with 500GB standard storage.

The pilot is considered a success and your CEO has managed to get the attention of some potential investors. The business plan requires a deployment of at least 100K sensors which needs to be supported by the backend. You also need to store sensor data for at least two years to be able to compare year over year Improvements.

To secure funding, you have to make sure that the platform meets these requirements and leaves room for further scaling.

Which setup will meet the requirements?

- A. Add an SQS queue to the ingestion layer to buffer writes to the RDS instance
- B. Ingest data into a DynamoDB table and move old data to a Redshift cluster
- C. Replace the RDS instance with a 6 node Redshift cluster with 96TB of storage
- D. Keep the current architecture but upgrade RDS storage to 3TB and 10K provisioned IOPS

Answer: C

Explanation:

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

Question: 4

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 a 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

Question: 5

A company is storing data on Amazon Simple Storage Service (S3). The company's security policy mandates that data is encrypted at rest.

Which of the following methods can achieve this? (Choose 3)

- A. Use Amazon S3 server-side encryption with AWS Key Management Service managed keys.
- B. Use Amazon S3 server-side encryption with customer-provided keys.
- C. Use Amazon S3 server-side encryption with EC2 key pair.
- D. Use Amazon S3 bucket policies to restrict access to the data at rest.
- E. Encrypt the data on the client-side before ingesting to Amazon S3 using their own master key.
- F. Use SSL to encrypt the data while in transit to Amazon S3.

Answer: A,B,E

Question: 6

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 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 idle. Once data is processed, change the storage class of the S3 objects to Reduced Redundancy Storage.
- B. 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 Reduced Redundancy Storage.
- C. 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.
- D. Use SNS to pass job messages use Cloud Watch alarms to terminate spot worker instances when they become idle. Once data is processed, change the storage class of the S3 object to Glacier.

Answer: C

Question: 7

You've been hired to enhance the overall security posture for a very large e-commerce site. They have a well architected multi-tier application running in a VPC that uses ELBs in front of both the web

and the app tier with static assets served directly from S3. They are using a combination of RDS and DynamoDB for their dynamic data and then archiving nightly into S3 for further processing with EMR. They are concerned because they found questionable log entries and suspect someone is attempting to gain unauthorized access.

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

- A. Recommend that they lease space at a DirectConnect partner location and establish a 1G DirectConnect connection to their VPC they would then establish Internet connectivity into their space, filter the traffic in hardware Web Application Firewall (WAF). And then pass the traffic through the DirectConnect connection into their application running in their VPC.
- B. Add previously identified hostile source IPs as an explicit INBOUND DENY NACL to the web tier subnet.
- C. Add a WAF tier by creating a new ELB and an AutoScaling group of EC2 Instances running a 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

Question: 8

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 will 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 Kinesis with Amazon Kinesis and save the results to DynamoDB.

Answer: B

Question: 9

You are designing Internet connectivity for your VPC. The Web servers must be available on the Internet.

The application must have a highly available architecture.

Which alternatives should you consider? (Choose 2)

- A. Configure a NAT instance in your VPC. Create a default route via the NAT instance and associate it with all subnets. Configure a DNS A record that points to the NAT instance public IP address.
- B. Configure a CloudFront distribution and configure the origin to point to the private IP addresses of your Web servers. Configure a Route53 CNAME record to your CloudFront distribution.
- C. Place all your web servers behind ELB. Configure a Route53 CNAME to point to the ELB DNS name.
- D. Assign EIPs to all web servers. Configure a Route53 record set with all EIPs, with health checks and DNS failover.
- E. Configure ELB with an EIP. Place all your Web servers behind ELB. Configure a Route53 A record that points to the EIP.

Answer: C,D

Question: 10

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. 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.
- D. Create your RDS instance separately and pass its DNS name to your's DB connection string as an environment variable. Alter its security group to allow access to it from hosts in your application subnets.

Answer: A

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.

<http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/AWSHowTo.RDS.html>