

KillTest

高品質 より良いサービス



練習問題

<https://www.killtest.jp>

1年間無料アップデート

Exam : **AWS Certified Solutions Architect-Professional**

Title : AWS Certified Solutions Architect-Professional

Version : DEMO

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 three.)

- 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: ACD

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.

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

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 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

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 three.)

- 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: ABE

Explanation:

Reference: <http://docs.aws.amazon.com/AmazonS3/latest/dev/UsingKMSEncryption.html>