CRYPTO 101: REAL-WORLD DEPLOYMENTS

5. AWS KEY MANAGEMENT

Alfred Menezes
cryptography101.ca

March 2025

Outline

1. AWS global infrastructure
2. Data centre security
3. Hardware Security Modules (HSMs)
4. AWS encryption and key management

Encryption in the cloud: AWS

- AWS: Amazon Web Services
- Global cloud market share (2024):
  - AWS 31%
  - Microsoft Azure 24%
  - Google Cloud 11%
- AWS 2024:
  - Revenue $107.6 billion, Profits $39.8 billion
  - 17% of Amazon revenues, 67% of Amazon profits

AWS services:

- Compute: EC2 (Elastic Compute Cloud), ….
- Storage: S3 (Simple Storage Service), …
- Database: Aurora (relational database), DynamoDB (NoSQL database), …
- Security: Identity and Access Management (IAM), Certificate Manager, CloudHSM, Key Management Service (KMS), …

AWS global infrastructure

- 36 Regions
- 114 Availability Zones
  (at least three in each region).
  - Each availability zone is comprised of one or more data centres, which are physically separated from each other.
- 700+ Edge Locations (data centres).
- "Tens of millions of servers", replaced on a five-year cycle.

aws.amazon.com/about-aws/global-infrastructure

(March 2025)

AWS security

- AWS customers include Netflix, Apple, Facebook, Twitch, Zoom, Reddit, US federal government (including Department of Defence, CIA), Canadian government, Coinbase, JP Morgan, Goldman Sachs, etc.
- Security is needed for: (a) data in transit, (b) data in use, (c) data at rest.
  - Data centre security
  - *Encrypt everything*
  - Hardware Security Modules (HSMs)

Current Efforts - Google

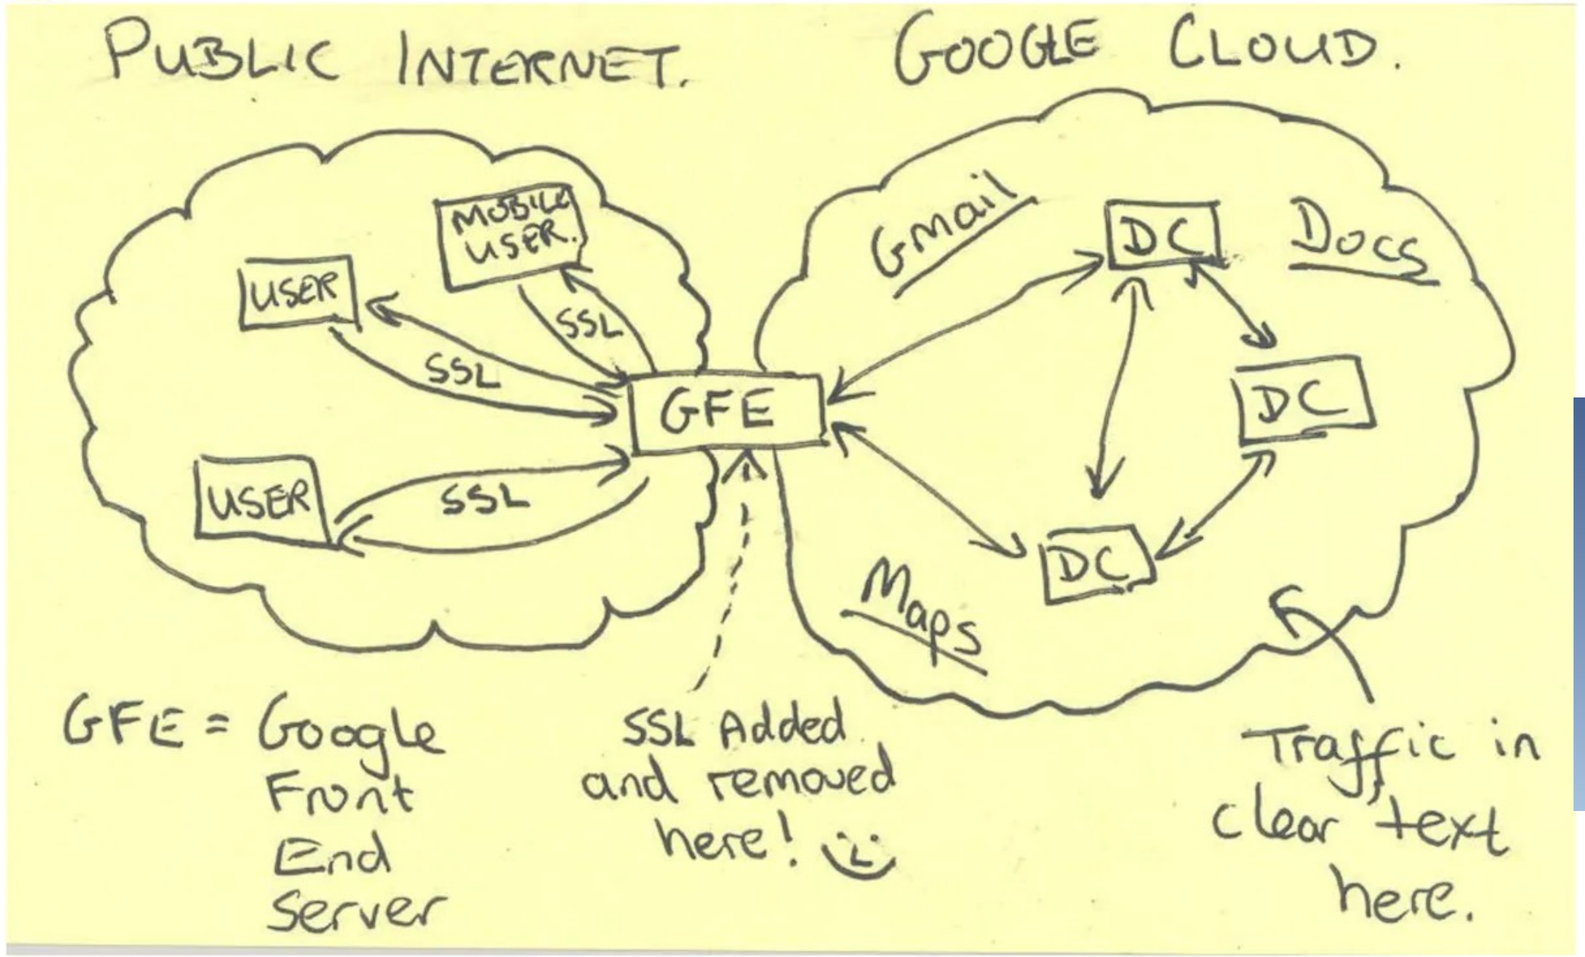

Snowden documents
October 30, 2013

An unnamed telecommunications operator provided GCHQ with secret access to its fibre optic cables that transported data between Google (and Yahoo!) data centres.

Data centre security

- Nondescript undisclosed facilities, with secured entrances.
- 24/7 monitoring: Professional security staff used video surveillance, intrusion detection, access log monitoring.
- Planning for flooding, extreme weather, seismic activity, fires, power supply, securely destroying old storage devices, etc.
- All data flowing between data centres is automatically encrypted.
- See: aws.amazon.com/compliance/data-center/data-centers/

Hardware Security Modules (HSMs)

- Expensive, tamper-resistant devices.
- Very limited API, e.g. generate key, encrypt/decrypt, delete key, generate random bytes.
- When operational, no AWS operator can access an HSM, e.g. via ssh, and no software updates are allowed.
- After reboot and in a non-operational state, there is no keying material on the HSM.
- Software can only be updated after multiple AWS employees have reviewed the code, and under quorum of multiple AWS operators with valid credentials.
- FIPS 140-2 Level-3 certified (an HSM certification program run by NIST and CSE). November 2023: tinyurl.com/aws-fips140

AWS encryption

- KMS: Key Management Service.
- Data at rest, for **all** data stored by **all** AWS services.
- Client-side/Server-side encryption, where the AWS KMS generates and manages keys.
- Encrypt everything: Each data item is encrypted with a ***one-time*** data encryption key (DEK).
  - Consequently, there are a *massive* number of encryption keys to manage (generation, storage, backup, retrieval, access control, deletion, etc.)

- This naturally leads to the notion of envelope encryption:
  - Each client has a Customer Main Key (CMK).
  - The CMK is used to encrypt/decrypt DEKs.
  - DEKs are used to encrypt/decrypt customer data.
  - DEKs are deleted after each usage.

Plaintext encryption

- $C = \text{Enc}_k(M)$, where Enc=AES-GCM.
- $IV \in \{0,1\}^{96}$, $J_0 = IV \| 0^{31}1$, so $u \leq 2^{32} - 2$.
- Plaintexts can be broken into chunks, and each chuck encrypted with a different IV.
- The IV can be a counter if it is easy to maintain state.
- Otherwise, the IV is selected at random.

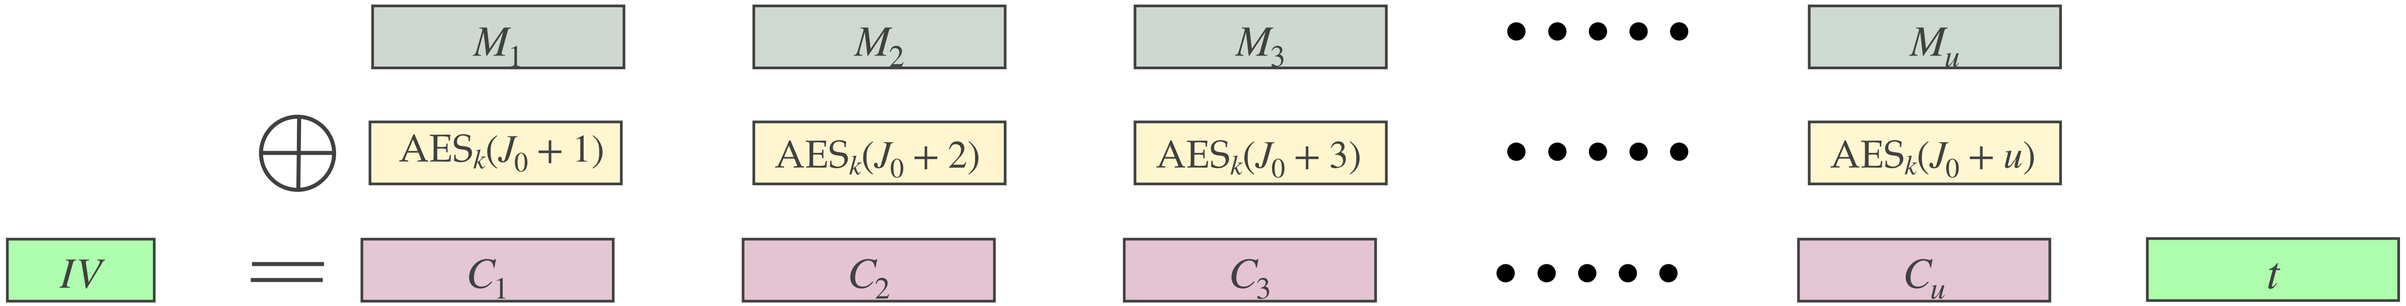

Random IVs

- NIST: The IVs in GCM must fulfill the following "uniqueness" requirement:
  - The probability that the authenticated encryption function ever will be invoked with the same IV and the same key on two (or more) distinct sets of input data shall be no greater than 2^{-32}.
- Birthday paradox: If B IV's are drawn at random (with replacement) from $\{0,1\}^{96}$, the probability that at least one string is selected two (or more) times is $\approx B^2/(2 \cdot 2^{96})$.
- So, at most $B = 2^{32}$ chunks should be encrypted with the same DEK k when IVs are selected at random.

DEK encryption

- $wk = \text{Enc}_K(A, k)$, where Enc=AES-GCM.
- $IV \in \{0,1\}^{96}$, $J_0 = IV \| 0^{31}1$.
- A CMK K can be used simultaneously by several HSMs. State management between these HSMs should be minimized to reduce latency. So, IVs are selected at random.
- NIST uniqueness requirement $\Longrightarrow$ K can be used to wrap at most 2^{32} DEKs.
- This can be problematic, since CMK's are rotated only once a year.
- So, only 136 DEKs can be encrypted per second (on average).

Derive key mode for key wrapping

- Matt Campagna & Shay Gueron.
- Select a random nonce $N \in_R \{0,1\}^{128}$.
  1. Derive a one-time 256-bit key $L = \text{HMAC}_K(N)$.
  2. Key wrap with L: $wk = \text{Enc}_L(A, k),\ N$.
- A careful analysis shows, for example, that the system meets the NIST requirement, even when 2^{40} CMKs are used, each encrypting as many as 2^{50} DEKs.

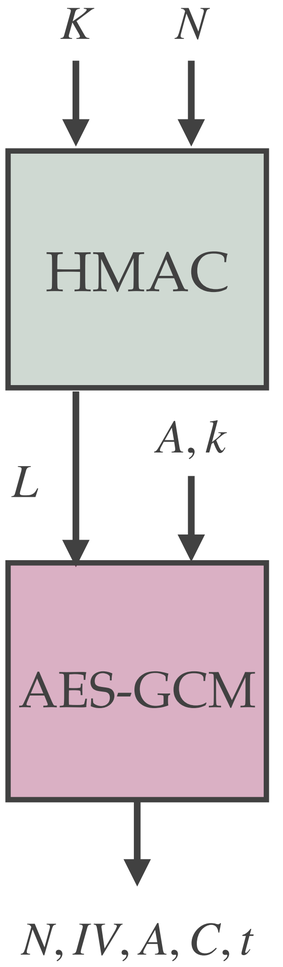

KMS pricing gives an idea of AWS KMS scale

- Quota on cryptographic operations request rate (encrypt, decrypt, etc.):
  - To ensure that AWS KMS can provide fast and reliable responses to API requests from all customers, it throttles API requests that exceed certain boundaries.
  - Quota on API requests, 10,000/second (20,000 and 100,000 in some regions) — can be increased by request.
- Pricing:
  - CMK creation: $1/month (prorated hourly)
  - KMS API requests per month: 20,000 free, $0.03/10,000 requests.
- S3 example:
  - 1 CMK, used to encrypt 10,000 files, that are collectively decrypted for access 2,000,000 times per month.
  - Cost: $1 + (10,000 + 2,000,000 - 20,000) * 0.03/10000 = $6.97/month.

DynamoDB

- Amazon's DynamoDB is a fully-managed proprietary NoSQL database.
- Offered as an AWS service.
- ***Provides single-digit millisecond latency at virtually any scale.***
- Each table entry is encrypted with a unique DEK.
- Handles about 90 million requests per second (as of 2022).
- Example: (2022) Snapchat stores all the Snap *metadata* on DynamoDB.
  - 400+ Terabytes of metadata (the data itself (2+ exabytes) is stored on S3).
  - Performs nightly scans: 2+ billion row scans per minute.
- DynamoDB encryption client:
  - Java and Python libraries.
  - Support for client-side encryption, whereby you encrypt data items before sending them to DynamoDB for storage in the AWS cloud.

Envelope encryption

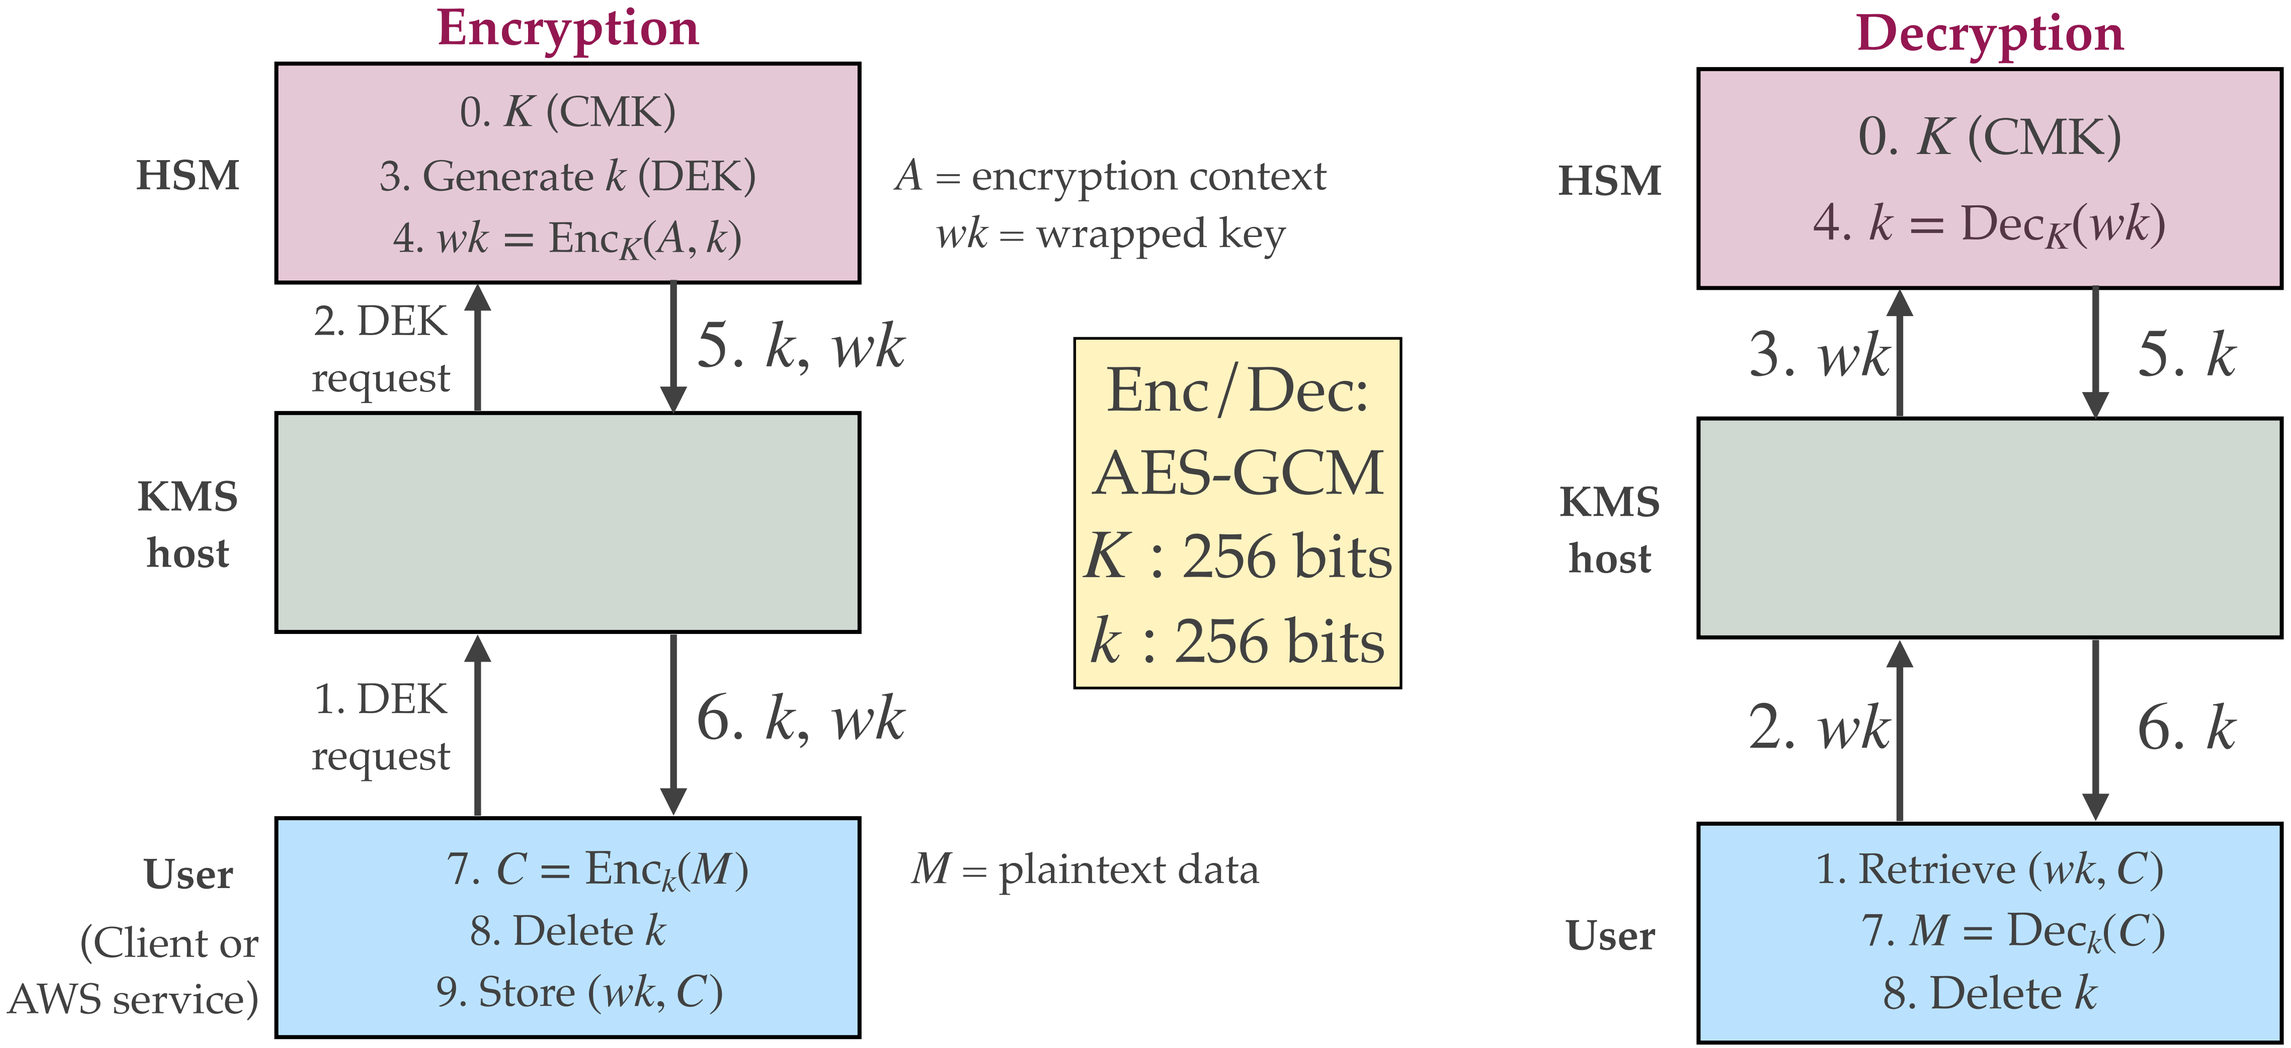

Protecting a CMK

- A CMK must have high security, availability, and durability.
- A CMK never leaves an HSM (in unencrypted form).
- Copies of a CMK in use are stored in memory (never on disk) of the HSMs in a domain — a collection of HSMs in a region, with at least three HSMs per availability zone. The HSMs in a domain do not communicate directly with each other.
- A CMK is encrypted with a domain key — a 256-bit AES-GCM key that is shared by all HSMs in the domain.
- The ciphertext, called an Exported Key Token, is stored in highly durable, low-latency storage.
- Exported key tokens are managed by the KMS.

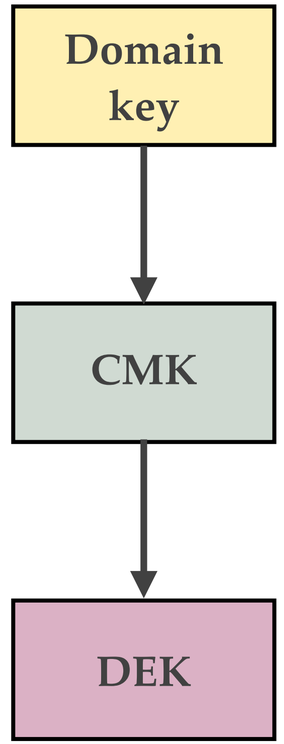

Exported Key Tokens (EKTs)

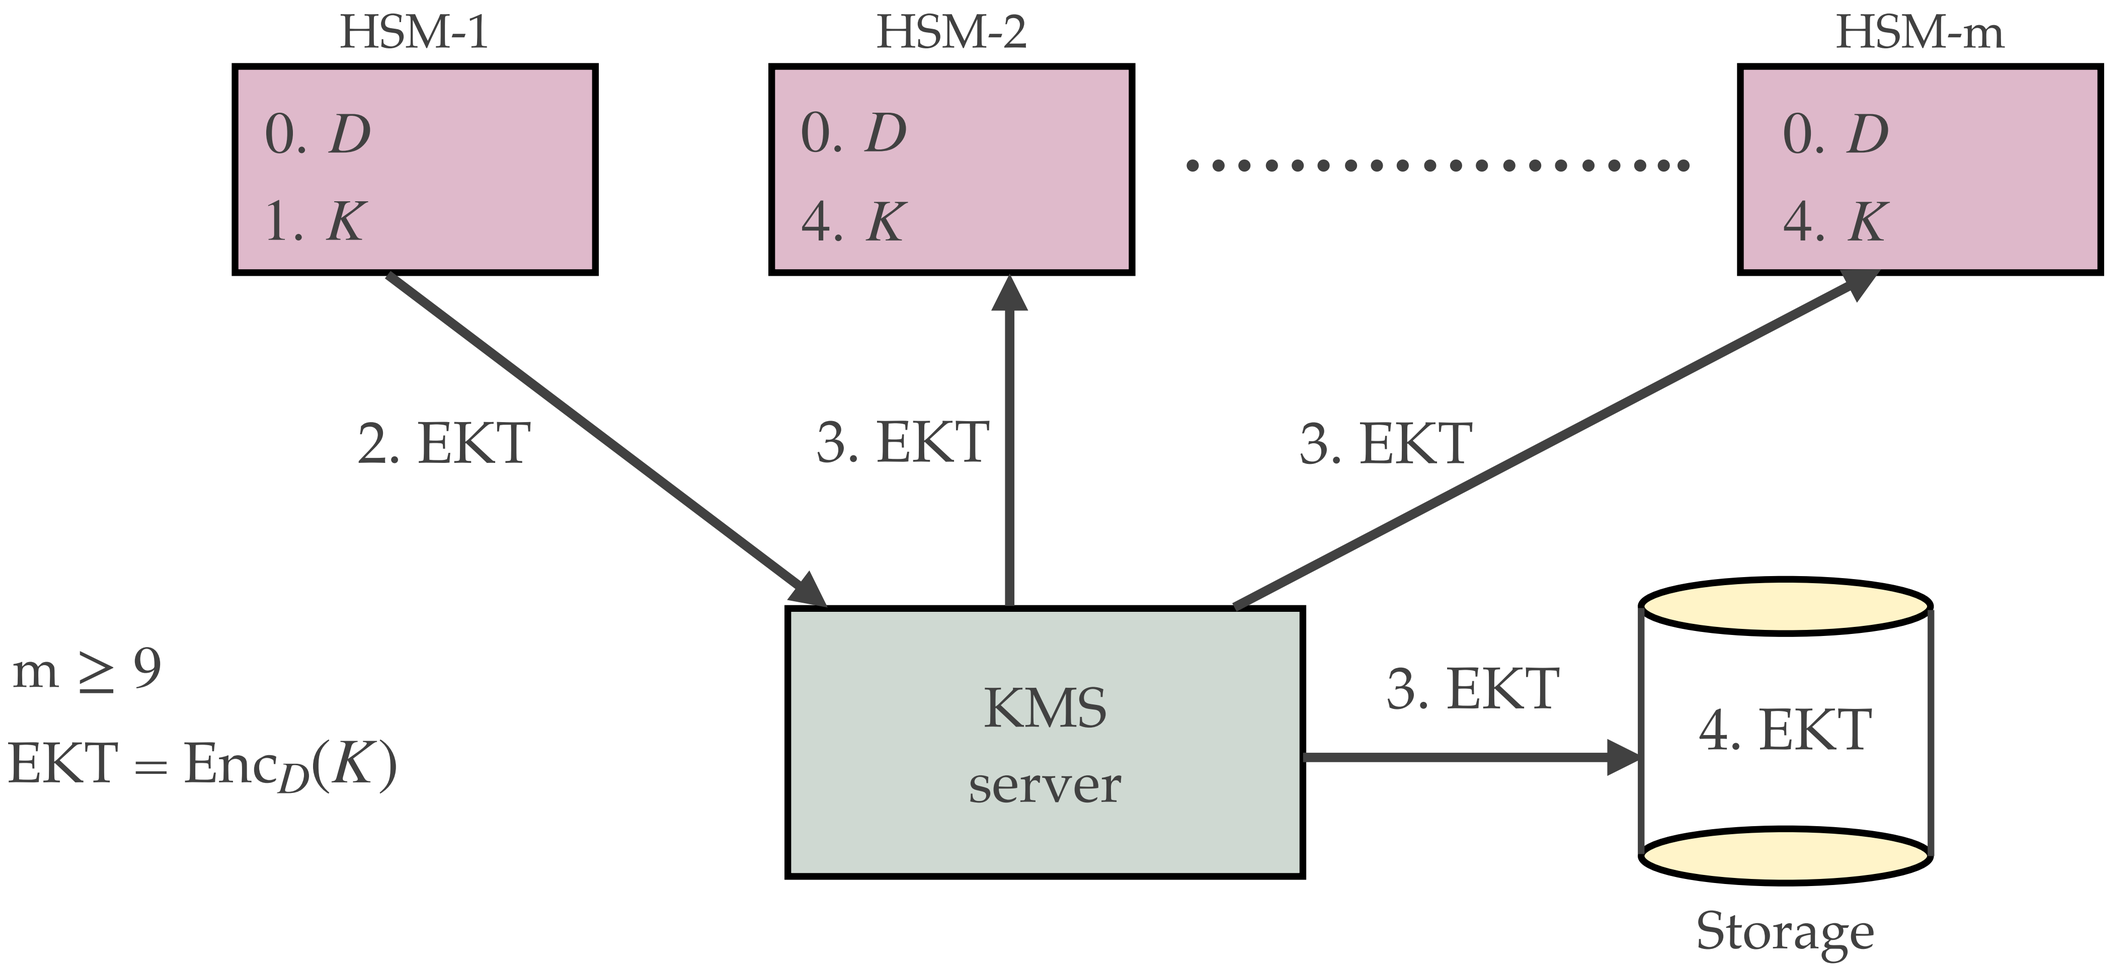

$m \geq 9$

$EKT = Enc_D(K)$

Protecting a domain key (AWS key hierarchy)

- A domain key is rotated daily.
  During rotation, encrypted CMKs are re-encrypted with the new domain key.
- One of the HSMs in a domain, say HSM-1, selects the new domain key, and encrypts it with a one-time AES-GCM key ℓ; let the resulting ciphertext be denoted c.
- The key ℓ is encrypted using public-key encryption (elliptic curve cryptography) — using the elliptic curve public keys of the other HSMs in the domain.
- Among the HSMs in a domain are some offline HSMs. The offline HSMs are stored in safes within monitored safe rooms in multiple independent geographical locations. Each safe requires at least one AWS security officer and one AWS KMS operator, from two independent AWS teams, to obtain these materials.
- Let the resulting ciphertexts be denoted $c_1, c_2, \ldots, c_n$.
- The list of ciphertexts $C = (c, c_1, c_2, \ldots, c_n)$ is signed by HSM-1.
- The signed C, called the exported domain token (EDT), is distributed by the KMS to the other (online) HSMs in the domain, and is also stored in highly durable storage.

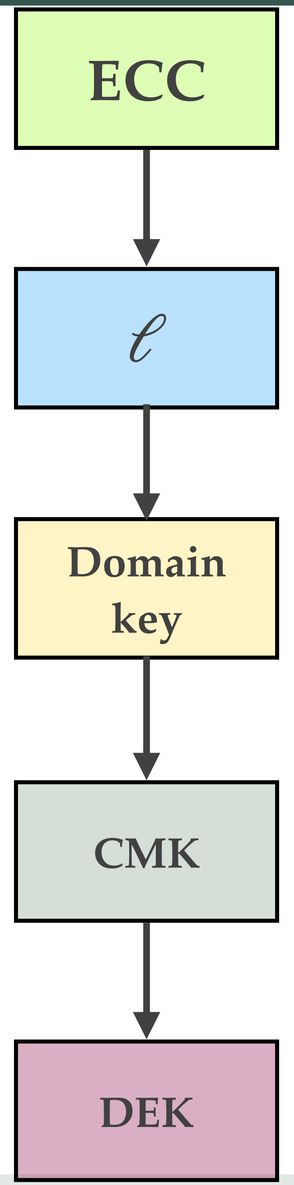

Exported Domain Tokens (EDTs)

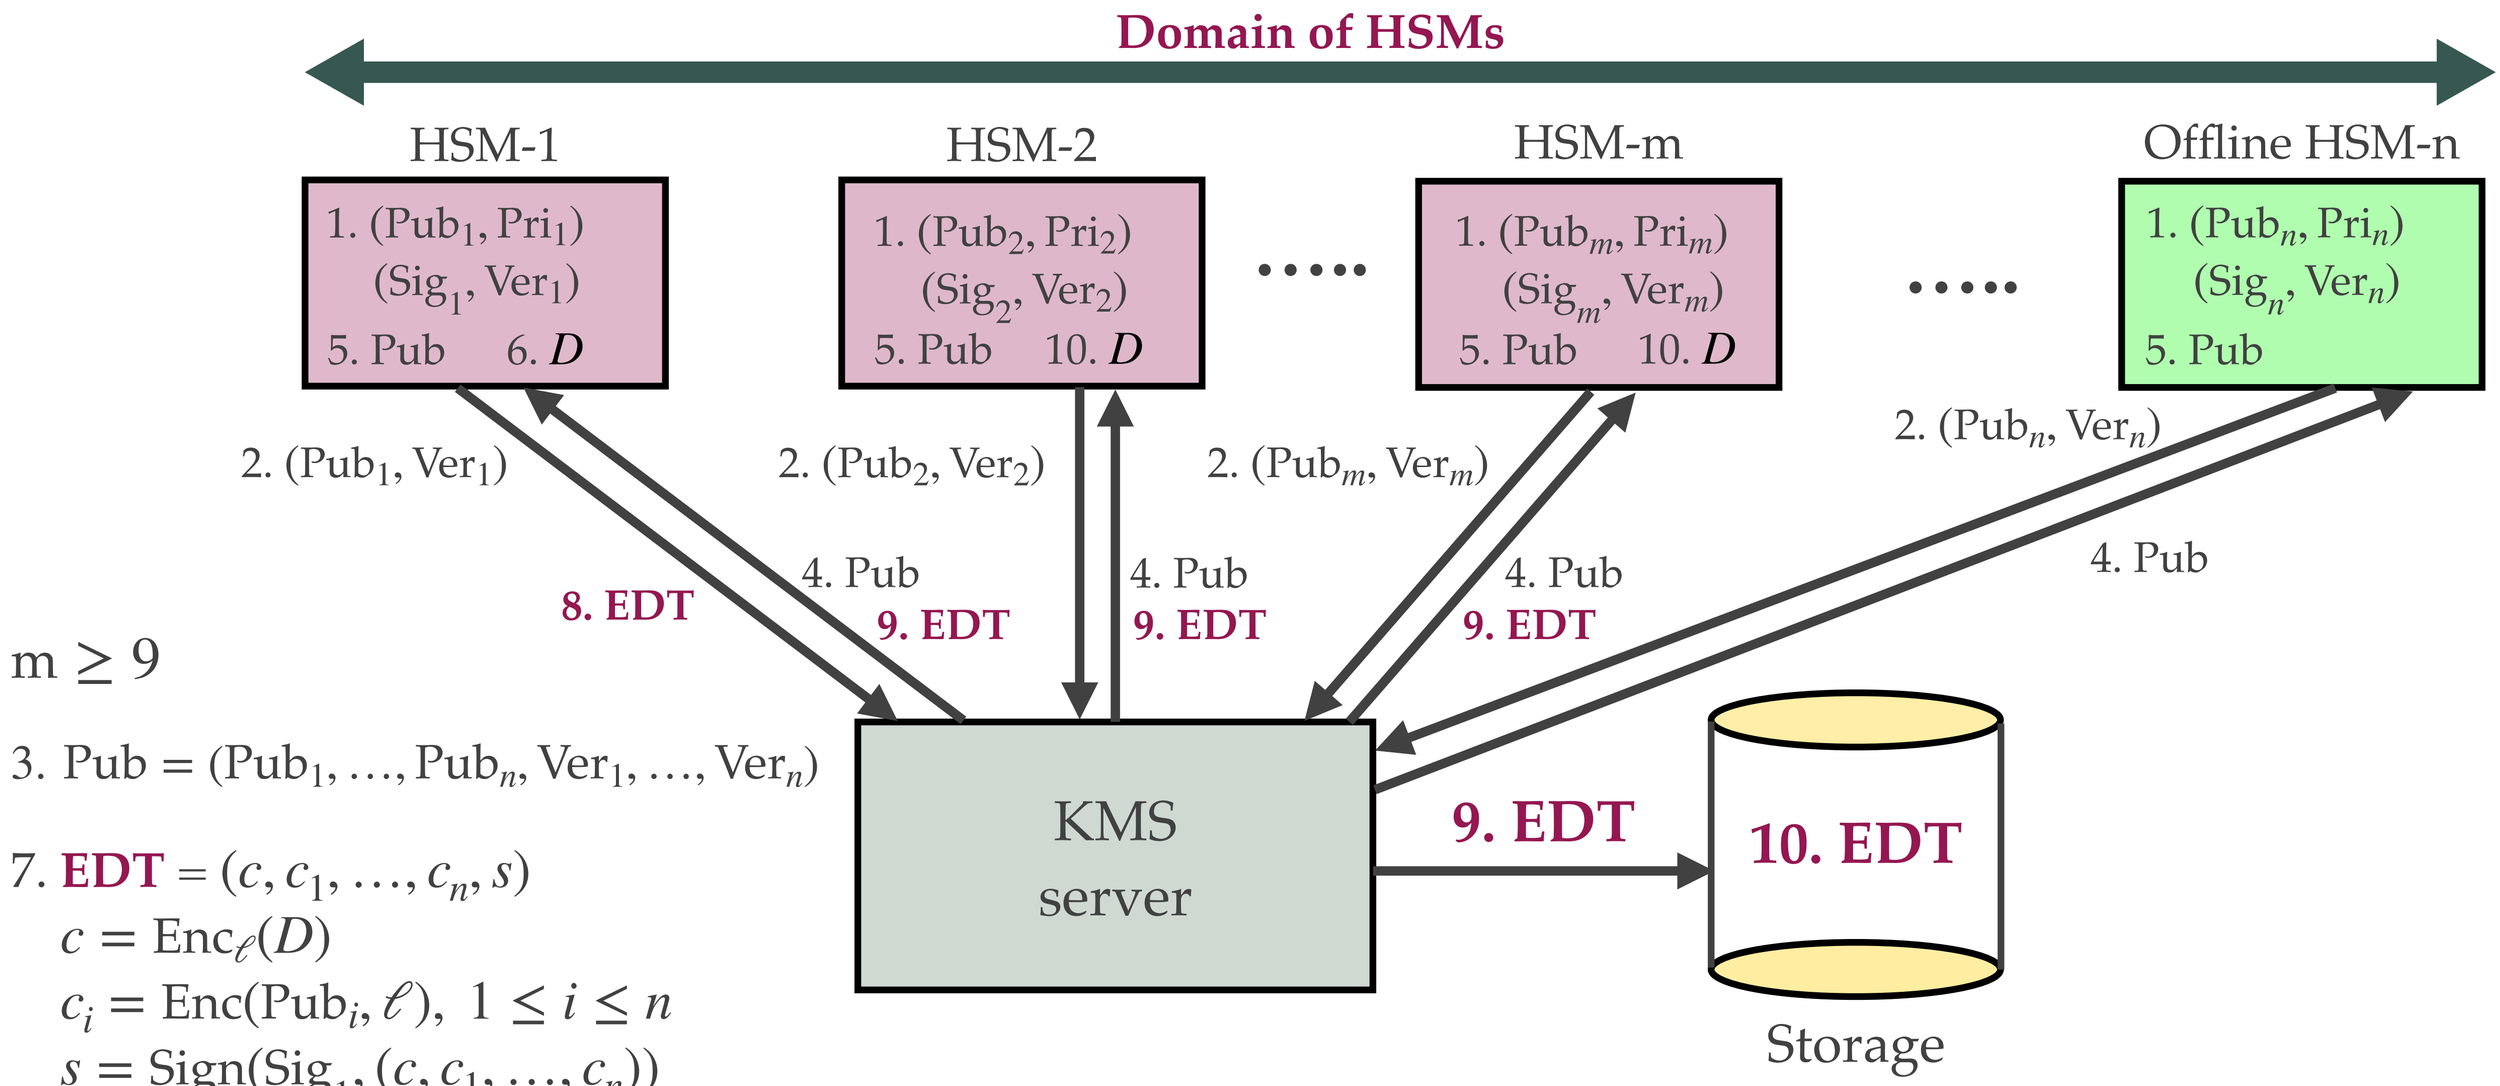

$m \geq 9$

3. $\text{Pub} = (\text{Pub}_1, \ldots, \text{Pub}_n, \text{Ver}_1, \ldots, \text{Ver}_n)$

7. **EDT** $= (c, c_1, \ldots, c_n, s)$
   $c = \text{Enc}_{\ell}(D)$
   $c_i = \text{Enc}(\text{Pub}_i, \ell), \ 1 \leq i \leq n$
   $s = \text{Sign}(\text{Sig}_1, (c, c_1, \ldots, c_n))$

AWS references

1. AWS KMS cryptographic details
   tinyurl.com/kms-details

2. AWS encryption SDK: Developer guide
   tinyurl.com/aws-sdk

3. M. Campagna and S. Gueron,
   "Key management systems at the cloud scale", *Cryptography*, 2019.