

Message integrity: Confidentiality alone not sufficient, also need message integrity. Otherwise adversary can tamper with the message
 (e.g., "Send \$100 to Bob" → "Send \$100 to Eve")

In some cases (e.g., software patches), integrity more important than confidentiality

Idea: Append a "tag" (also called a "signature") to the message to prove integrity (property we want is tags should be hard to forge)

Observation: The tag should be computed using a keyed-function

↳ Example of keyless integrity check: CRC (cyclic redundancy check) [simple example is to set tag to be the parity] ^{better}

↳ this was used in SSH v1 (1995) for data integrity! Fixed in SSH v2 (1996)

↳ also used in WEP (802-11b) protocol for integrity - also broken!

this tolerates a single error

(better error-correcting codes can do much better)

Problem: If there is no key, anyone can compute it! Adversary can tamper with message and compute the new tag.

Definition. A message authentication code (MAC) with key-space K , message space M and tag space T is a tuple of algorithms $\Pi_{MAC} = (\text{Sign}, \text{Verify})$:

$$\text{Sign}: K \times M \rightarrow T$$

$$\text{Verify}: K \times M \times T \rightarrow \{0,1\}$$

} Must be efficiently-computable

Correctness: $\forall k \in K, \forall m \in M$:


$$\Pr[\text{Verify}(k, m, \text{Sign}(k, m)) = 1] = 1$$

↳ Sign can be a randomized algorithm

Defining security: Intuitively, adversary should not be able to compute a tag on any message without knowledge of the key

↳ Moreover, since adversary might be able to see tags on existing messages (e.g., signed software updates), it should not help towards creating a new MAC

Definition. A MAC $\Pi_{MAC} = (\text{Sign}, \text{Verify})$ satisfies existential unforgeability against chosen message attacks (EUF-CMA) if for all efficient adversaries A , $\text{MAC}_{adv}[A, \Pi_{MAC}] = \Pr[W=1] = \text{negl}(\lambda)$, where W is the output of the following security game:



As usual, λ denotes the length of the MAC secret key (e.g., $\log |K| = \text{poly}(\lambda)$)

Note: the key can also be sampled by a special KeyGen algorithm (for simplicity, we just define it to be uniformly random)

Let m_1, \dots, m_Q be the signing queries the adversary submits to the challenger, and let $t_i \leftarrow \text{Sign}(k, m_i)$ be the challenger's responses. Then, $W=1$ if and only if:

$$\text{Verify}(k, m^*, t^*) = 1 \text{ and } (m^*, t^*) \notin \{(m_1, t_1), \dots, (m_Q, t_Q)\}$$

MAC security notion says that adversary cannot produce a new tag on any message even if it gets to obtain tags on messages of its choosing.

First, we show that we can directly construct a MAC from any PRF.

MACs from PRFs: Let $F: K \times M \rightarrow T$ be a PRF. We construct a MAC Π_{MAC} over (K, M, T) as follows:

Sign(k, m): Output $t \leftarrow F(k, m)$

Verify(k, m, t): output 1 if $t = F(k, m)$ and 0 otherwise

Theorem. If F is a secure PRF with a sufficiently large range, then Π_{MAC} defined above is a secure MAC. Specifically, for every efficient MAC adversary A , there exists an efficient PRF adversary B such that


$$\text{MACAdv}[A, \Pi_{\text{MAC}}] \leq \text{PRFAdv}[B, F] + \frac{1}{|T|}.$$

Intuition for proof: 1. Output of PRF is computationally indistinguishable from that of a truly random function.

2. If we replace the PRF with a truly random function, adversary wins the MAC game only if it correctly predicts the random function at a new point. Success probability is then exactly $1/|T|$.

Proof. We define the following sequence of hybrid experiments:

Hyb₀: This is the MAC security game:




Goal: Show for all efficient A :
 $\Pr[\text{Hyb}_0(A) = 1] = \text{negl.}$

Experiment outputs 1 if adversary did not query on m^* and $t^* = F(k, m^*)$

Hyb₁: Same as Hyb₀ except we replace $F(k, \cdot)$ with $f(\cdot)$ where $f \xleftarrow{\$} \text{Funs}[M, T]$

Lemma 1. If F is a secure PRF, then for all efficient adversaries A ,
 $|\Pr[\text{Hyb}_0(A) = 1] - \Pr[\text{Hyb}_1(A) = 1]| = \text{negl.}$

Proof. Suppose there exists efficient A such that above probability is ϵ . We construct B as follows:



$$\left. \begin{array}{l} \Pr[B \text{ outputs } 1 \mid b=0] = \Pr[\text{Hyb}_0(A) = 1] \\ \Pr[B \text{ outputs } 1 \mid b=1] = \Pr[\text{Hyb}_1(A) = 1] \end{array} \right\} \text{PRFAdv}[B, F] = \epsilon$$

Lemma 2. For all adversaries A , $\Pr[\text{Hyb}_1(A) = 1] = \frac{1}{|T|}$.

Proof. Hyb₁(A) outputs 1 if A predicts value of f at m^* . Since f is uniform, A succeeds with probability at most $1/|T|$.

Implication: Any PRF with large output space can be used as a MAC.


↳ AES has 128-bit output space, so can be used as a MAC

Drawback: Domain of AES is 128-bits, so can only sign 128-bit (16-byte) messages

How do we sign longer messages? We will look at two types of constructions:

1. Constructing a large-domain PRF from a small-domain PRF (i.e., AES)
2. Hash-based constructions

Approach 1: use CBC (without IV)



Not encrypting messages so no need for IV (or intermediate blocks)

↳ Mode often called "raw-CBC"



Raw-CBC is a way to build a large-domain PRF from a small-domain one

↳ Can show security for "prefix-free" messages [more precisely, raw-CBC is a prefix-free PRF: pseudorandom as long as PRF never evaluated on two values where one is a prefix of other]

↳ includes fixed-length messages as a special case

But not secure for variable-length messages: "Extension attack"

1. Query for MAC on arbitrary block x :



2. Output forgery on message $(x, x \oplus t)$ and tag t

$\Rightarrow t$ is a valid tag on extended message $(x, t \oplus x)$

↳ Adversary succeed with advantage 1


raw CBC can be used to build a MAC on fixed-length messages, but not variable-length messages (more generally, prefix-free)

For variable-length messages, we use "encrypted CBC": ^(ECBC) standards for banking / financial services

↳ variant used in ANSI X9.9, ANSI X9.19 standards

critical for security (using the same key not secure)

apply another PRF with a different key to the output of rawCBC



To use encrypted CBC-MAC, we need to assume message length is even multiple of block size (similar to CBC encryption)

↳ to sign messages that are not a multiple of the block size, we need to first pad the message

↳ as was the case with encryption, padding must be injective

↳ in the case of encryption, injectivity needed for correctness

↳ in the case of integrity, injectivity needed for security [if $\text{pad}(m_0) = \text{pad}(m_1)$, m_0 and m_1 will have the same tag]

Standard approach to pad: append $1000\dots 0$ to fill up block [ANSI X9.9 and ANSI X9.19 standards]

- Note: if message is an even multiple of the block length, need to introduce a dummy block

↳ Necessary for any injective function: $| \{0,1\}^{2^n} | > | \{0,1\}^n |$

- This is a bit-padding scheme [PKCS #7 that we discuss previously in the context of CBC encryption is a byte-padding scheme]

Encrypted CBC-MAC drawbacks: always need at least 2 PRF evaluations (using different keys) } especially bad for authenticating short (e.g., single-byte) messages
 messages must be padded to block size

Better approach: raw CBC-MAC secure for prefix-free messages

↳ Can we apply a "prefix-free" encoding to the message?

equal-length messages cannot have one be prefix of other
 different-length messages differ in first block

- Option 1: Prepend the message length to the message

Problematic if we do not know message length at the beginning (e.g., in a streaming setting)


Still requires padding message to multiple of block size

- Option 2: Apply a random secret shift to the last block of the message

$$(x_1, x_2, \dots, x_L) \mapsto (x_1, x_2, \dots, x_L \oplus k) \text{ where } k \in \mathbb{F}_X$$

Adversary that does not know k cannot construct two messages that are prefixes except with probability $1/|X|$ (by guessing k)

Cipher-based MAC (CMAC): variant of CBC-MAC standardized by NIST in 2005 } randomized prefix-free encoding
 clever technique to avoid extra padding block
 better than encrypted CBC (should be preferred over ANSI standards)



secret random shift (part of the MAC key)

different keys needed to avoid collision between unpadded message and padded message
 ending in $100\dots 0$

if message is not a multiple of block length, then pad (ANSI) and xor with different secret key k_2

never needs to introduce an additional block!

key: (k, k_1, k_2) - CMAC standard uses a specific key-derivation function to derive these keys from one key