Cryptographic Engineering

working draft

Matthew D. Green
Johns Hopkins University

Excerpt: Chapter 5, Symmetric encryption and authentication

§5.3.6 Nonces and misuse-resistance (book pages 91-91)

Posted for 601.445/601.645 Practical Cryptographic Systems, Fall 2026
Lecture 19, Wed Nov 4: implementation failures (nonce reuse)

`https://practicalcrypto.org`

This is an excerpt from an unfinished book draft: expect gaps, rough edges, and errors, and note that section numbers and page numbers may change in later drafts. Cross-references to other chapters point to material that is not included here. Copyright © 2026 Matthew D. Green. All rights reserved. Posted for the students of the course above; **please do not redistribute** — link to the course site instead. Built 2026-10-19.

5.3.6 Nonces and misuse-resistance

AEAD schemes require the encryptor to select a nonce prior to encrypting each message. As mentioned above, the same nonce should never be used twice with a given key. However, mistakes do happen! Since implementation errors are common, cryptographers have sought to minimize the damage caused by such a bug. The results are called *misuse-resistant* schemes [284, 150, 23].

The main goal of a misuse-resistant AEAD scheme is to prevent catastrophic failure in the event that nonces are re-used to encrypt messages. AEAD schemes can be classified into three categories:

1. **Not misuse-resistant.** Here, the re-use of a small number of nonces can have catastrophic effects on the security of all messages encrypted with that key. For example, both of the popular GCM and ChaCha-Poly1305 schemes leak the plaintext of messages encrypted with the same nonce; however (and even worse), they will leak the scheme's internal authentication key if a small number of nonces are re-used. This exposes the whole scheme to tampering attacks [184].

2. **Partially misuse-resistant.** If nonces are re-used, confidentiality may be lost entirely for the specific messages where nonce re-use occurs. However, all other ciphertexts where nonces are *not* repeated will remain fully secure, and authenticity is maintained. CCM mode is an example of a mode that provides this property [168].

3. **Fully misuse-resistant.** When the same nonce is re-used to encrypt two messages M_1, M_2 the most information that an attacker will learn is whether $M_1 = M_2$. All other security properties of the scheme remain, including full confidentiality for M_1, M_2 when $M_1 \neq M_2$. An example of a scheme with this property is GCM-SIV [149, 150].

DRAFT

The results in the "partial" category typically stem from the use of a stream cipher to encrypt data. Here the re-use of a nonce to encrypt messages M_1, M_2 will produce ciphertexts of the form $M_1 \oplus KS$, $M_2 \oplus KS$ where KS is the specific keystream produced for that nonce. An attacker can XOR these values to obtain $M_1 \oplus M_2$, which may leak information about either message.

Synthetic IVs. Fully misuse-resistant schemes can avoid this problem by constructing a second "synthetic" IV that's a function of the key, plaintext and nonce (as well as possibly the associated data.) The benefit is that even if the caller repeats a nonce, the synthetic IV should only repeat if *exactly* the same data is encrypted (including the associated data), which means that the attacker does not get the benefit of seeing different messages (or hash values) encrypted using the same IV. Schemes like GCM-SIV [149, 150] compute this IV very efficiently. Unfortunately there's at least some cost to this safety benefit: computing the synthetic IV increases the cost of encrypting a message.