

Analysis of VAES3 (FF2)

Morris Dworkin Ray Perlner

National Institute of Standards and Technology
Gaithersburg, Maryland, USA

April 2, 2015

Abstract

This note describes a theoretical chosen-plaintext attack on the VAES3 mode for format-preserving encryption. VAES3 was specified under the name FF2 in Draft NIST Special Publication 800-38G.

1 Summary

The National Institute of Standards and Technology (NIST) specified three methods for format-preserving encryption (FPE) in Draft NIST Special Publication (SP) 800-38G, which was released for public comment in July, 2013. Each method was a mode of operation of the Advanced Encryption Standard (AES) [1]. One of the three modes, VAES3—submitted to NIST in [2]—was specified under the name FF2 in the NIST draft. This note describes a theoretical chosen-plaintext attack that shows the security strength of FF2 is less than 128 bits.

2 Description of VAES3

VAES3 encrypts the plaintext, P , under a master key, K , and tweak, T , in two stages: 1) a 128-bit subkey, SK , is generated from an encoding of the tweak by invoking the AES block cipher under the master key; and 2) the subkey is used as the key for a Feistel-based encryption, FEIST, of the plaintext. Symbolically,

$$SK = \text{AES}_K(\text{encoding of } T)$$

$$\text{VAES3}_K(P, T) = \text{FEIST}_{SK}(P).$$

The Feistel-based encryption includes ten invocations of the AES block cipher as the pseudorandom round function, but this fact and the other details of FEIST do not affect the analysis below. The encoding of T includes the lengths of the plaintext and the tweak.

3 Chosen-Plaintext Attack

VAES3 invokes FEIST with a different subkey for every tweak. Therefore, many subkeys may be available as targets for an attack; moreover, given the right kind of information, the computational cost of a brute-force search is reduced.

In particular, suppose an attacker knew a set of ciphertexts $C_i = \text{VAES3}_K(P, T_i)$, where the set is indexed by i ; in other words, the inputs to VAES3 are “chosen” in the sense that distinct tweaks are paired to a single common plaintext, P . The actual values of P and the T_i may be arbitrary.

The attacker computes $D_j = \text{FEIST}_{K_j}(P)$ for some set of distinct subkeys, K_j , which also may be arbitrarily chosen. If $C_i = D_j$ for any pair of indices (i, j) , then K_j is a candidate subkey for the attacker to investigate further. If K_j turns out to be a false positive, then the attacker searches for another candidate. However, if confirmed, the knowledge that $K_j = \text{AES}_K(\text{encoding of } T_i)$ constitutes a significant breach of VAES3: for any other plaintext with the same length and with the same tweak T_i , the attacker could encrypt the plaintext, or decrypt its ciphertext, without discovering the master key.

In order to investigate a candidate K_j , the attacker must learn/know $\text{VAES3}_K(Q, T_i)$ for some number of additional of plaintexts Q with the same length (and format) as P . In particular, if $\text{VAES3}_K(Q, T_i) \neq \text{FEIST}_{K_j}(Q)$ for any Q , then K_j is a false positive. The number of successful trials that are necessary to confirm K_j as a true match depends on the length and format of P : if P contains b bits of information, then approximately $\lceil 128/b \rceil$ successful trials are required.

4 Estimate of the Complexity of the Attack

The probability that the attacker would find a true match depends on 1) the number of chosen inputs available and 2) the number of Feistel encryptions that the attacker can generate of the common plaintext with different subkeys. In particular, if there are 2^u and 2^v of them, respectively, for positive integers u and v , then at least one match would be expected when $u + v \geq 128$. Consequently, the security level of VAES3 should be regarded as no higher than $128 - u$.

Although not really practical, this level of computation is at least conceivable for an important intended use of FPE, the encryption of the middle-six digits of credit card numbers, tweaked by the outer ten digits. Thus, a single plaintext could in principle have 10^{10} different tweaks, so the size of the chosen input set could approach 2^{34} . In this case, the analysis would have a good chance of revealing a subkey if the attacker could compute FEIST for up to 2^{94} candidate subkeys. Moreover, multiple tweaks could potentially be compromised by the same set of chosen inputs.

Note that the computational costs of the attack are independent of the length of the master key. In other words, 192-bit or 256-bit master keys offer no additional security against this attack compared to 128-bit keys.

5 Acknowledgements

The authors acknowledge the National Security Agency for notifying NIST in general terms that FF2 might not meet NIST's security requirements.

References

- [1] FIPS Publication 197, *The Advanced Encryption Standard (AES)*, U.S. DoC/NIST, November 26, 2001, available at <http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf>.
- [2] Joachim Vance, *VAES3 scheme for FFX: An addendum to "The FFX Mode of Operation for Format-Preserving Encryption"*, Draft 1.0, May 20, 2011, available at <http://csrc.nist.gov/groups/ST/toolkit/BCM/documents/proposedmodes/ffx/ffx-ad-VAES3.pdf>.