Status Update of ForkAE

Elena Andreeva[1], Virginie Lallemand[2], Antoon Purnal[3], Reza Reyhanitabar[4], Arnab Roy[5], and Damian Vizár[6]

[1]DTU, Denmark
[2]Université de Lorraine, CNRS, Inria, LORIA, France
[3]imec-COSIC, KU Leuven, Belgium
[4]Elektrobit Automotive GmbH, Germany
[5]University of Klagenfurt, Austria
[6]CSEM, Switzerland

September 18, 2020

Abstract

The present document is the "Status Update" document for the second round NIST LWC candidate ForkAE, as called for by NIST LWC team mid-august 2020, summarizing (i) new security proofs for SAEF mode, (ii) third party analyses of the ForkSkinny primitive, (iii) third cryptanalysis of the directly related SKINNY tweakable blockcipher, (iv) planned tweaks for the final round, (v) new application use cases for the ForkAE family, as well as (vi) the new implementation aspects of the submission.

1 New proofs/arguments supporting security claims

1.1 OAE security of SAEF

In a recent work Andreeva, Singh Bhati and Vizár (under ongoing review process) show that the SAEF mode of operation achieves Online AE (OAE) security. OAE security notion is strictly stronger than the basic nonce-based AEAD security guarantees originally claimed for SAEF.

Online AE Security. Online Authenticated Encryption (OAE) by Fleischmann, Forler and Lucks (FSE'12) (and corrected by Hoang, Reyhanitabar, Rogaway, and Vizár in CRYPTO'15) captures a level of security between the basic nonce-based AEAD security and the full-fledged MRAE by Rogaway and Shrimpton (EUROCRYPT'06), achievable by online AE schemes. An online AE scheme processes the plaintext on the fly, such that i-th ciphertext block can be computed after having seen the first i plaintexts, allowing for a constant memory footprint in practice.

An OAE-secure scheme will leak the length of longest block-aligned prefix of two plaintexts encrypted with the same nonce and associated data but nothing more. As showed later by Endignoux and Vizár (FSE'17), OAE schemes are also

resistant to block-wise adaptive attacks (where an application outputs a part of the ciphertext before it has been fed the entire plaintext).

OAE Security of SAEF. SAEF has been proven to be OAE-secure as long as the total number of blocks processed with the same key is $\ll 2^{n/2}$ with n the blocksize of the underlying forkcipher ($n = 128$ for all SAEF instances in the ForkAE family). Such a birthday-bounded security is common among OAE schemes (COLM is also birthday secure). SAEF thus provides guarantees that are strictly stronger qualitatively and unchanged quantitatively. To our knowledge, there are only four other NIST LWC second round candidates with provable claims on nonce misuse security.

2 Third-party Analysis of ForkSkinny and ForkAE

ForkSkinny On the Dec. 10, 2019, we launched a cryptanalysis challenge for ForkSkinny, detailed on our website: `https://www.esat.kuleuven.be/cosic/forkae/home/forkskinny-challenge/` to encourage public cryptanalysis. Till this date, to the best of our knowledge, there have been no cryptanalysis results that break ForkSkinny with the specified parameters in the challenge.

In an article published at ToSC 2020, Bariant, David and Leurent showed that the best attacks on SKINNY can be extended to one extra round for most ForkSkinny variants, and up to 3 rounds for ForkSkinny-128-256. These results do not contradict our security claims, and in particular do not threaten our proposal. The ForkSkinny continues to benefit from the comfortable security margin of SKINNY (for instance the best current cryptanalysis results for SKINNY-128-256 reach about half of the total number of rounds). On the contrary, we believe that these rather limited improvement are a good sign regarding the security of our proposal.

ForkAE On the 12th of February, Patrick Derbez, Willi Meier, Ling Song, Reinhard Lueftenegger, Lenka Marekova and Danping Shi posted a comment on the lwc-forum mailing list. The authors reported that '*if an application allows to use ForkAE with various nonce lengths it is possible to mount a forgery*'. We stress that their analysis is incompatible with our cryptographic API, since the nonce length is explicitly specified as a parameter of PAEF, which is chosen and fixed for an instance.

On the 26th of February, Huicong Liang, Hongjun Wu and Meiqin Wang reported a forgery attack on the mode of operation PAEF to us. We pointed and communicated the invalidity of their analysis and they agreed with us.

3 Comments on the Security of Forkcipher

An idea which was used in an earlier forkcipher instantiation - ForkAES, is the exploitation of the reconstruction query in a forkcipher. It is unlikely that attacks that require inverse ForkSkinny queries in the "reconstruction" direction (i.e. adversary submits the "left" ciphertext block and expects to get the value

of the corresponding "right" ciphertext block) are applicable to the algorithms in the ForkAE family. Even though the reconstruction of the ciphertext blocks chosen by the adversary is computed in the decryption algorithms *internally*, the adversary cannot learn the values of the reconstructed "right" ciphertext block in a blackbox attack, unless a forgery occurs (at which point the guarantees of ForkSkinny become become void anyway).

4 Related Cryptanalysis Results

Our proposal being based on SKINNY and any new cryptanalysis of SKINNY can potentially have an impact on the security of ForkSkinny. Thus, we give here a brief overview of the most important recent advances on its analysis.

Although there are many cryptanalysis results published on SKINNY, none of them breaks the full round SKINNY. The advancement made since the announcement of round 2 candidates is rather limited, and as of today the cipher still has a security margin close to half its total number of rounds. The main recent advances are the following:

In a DCC article, Zhao and coauthors proposed a new technique that allows to attack 28 rounds of SKINNY-128-384 with a related-tweakey rectangle attack, which is one more round than before.

At CSCML 2020, Dunkelman et al. revisited some attacks on SKINNY, identified issues with several of them (for which the patch requires to reduce the number of attacked rounds) and showed that the diffusion of the cipher is low, as biases after 8 full rounds of SKINNY can be observed. The maximum number of attacked rounds is not increased.

At FSE 2019, Song et al. re-evaluated the probability of previous boomerang distinguishers with the BCT technique and its extension. The probabilities are shown to be higher than previously evaluated. At the same conference, Ankele et al. developed new zero-correlation and integral attacks that rely on the linear tweakey schedule of SKINNY. In these two papers, the total number of attacked rounds is not improved.

In another ToSC paper, Zhang, Cao, Guo, and Pasalic proposed new techniques to find integral, truncated and impossible differential distinguishers. The authors claim the first attack reaching 16 rounds out of 40 of SKINNY-128-128 in the single-key model.

At ICICS 2019, Chen et al. improved the complexity of the previous MITM attack by considering the key-bridging technique.

5 Planned tweak proposals

The SKINNY tweakable block cipher comes with large security margin (of around 50%), a fact supported by a large body of cryptanalytic results.

We plan to propose a reduction of the number of rounds which are iterated in ForkSkinny. More specifically, for the NIST compliant version ForkSkinny-128-288, we plan to reduce the number of rounds at least by 5 rounds after

forking, i.e. $r_0 = r_1 = 26$ compared to the exisiting $r_0 = r_1 = 31$. We believe further reduction in the number of rounds for $r_{\texttt{init}}$ and r_0, r_1 are possible and we are currently investigating this. The feasibility of reducing of number of rounds further (for all versions of ForkSkinny) is also supported by the finding of Bariant et al. that the best attack on SKINNY can be extended by 1 round (for the parameters of our ForkSkinny instances).

We additionally intend to extend and complement the ForkAE family with instances of the RPAEF mode, presented at ASIACRYPT'19. RPAEF is similar to PAEF but optimized to efficiently handle longer queries (requiring just one instance of ForkSkinny with a longer tweak compared to PAEF). Thanks to this inclusion, ForkAE family will become more versatile; it will still be the go-to solution for shortest queries, but it will also have a solution for the longer ones.

Finally, our yet unpublished results (under ongoing review process) show that ForkSkinny lends itself to the construction of CTR-like modes that achieve BBB security and improve efficiency of encryption-only for messages of any size. This is an additional feature of the ForkAE family which we may include in the next round.

6 New Use Cases

The new analysis referenced in Section 1.1 applies to SAEF without any need for modification, implying that ForkAE submission includes algorithms that cover the intersection between "lightweight" and "defense-in-depth" use cases, with the latter being a new use cases reported for ForkAE. We further illustrate the relevance of the use case below

Nonce misuse in lightweight applications. Many lightweight applications cannot use the most robust HW (due to cost constraints), while being exposed to physical attackers, making it easy for attackers to artificially amplify sources of accidental nonce misuse.

Blockwise encryption for external flash. When embedded platforms receive a large amount of data (such as a new firmware image), this has to be stored (at least temporarily) in the cheap but vulnerable external flash memory, simply because the data does not fit anywhere else. When encrypting such data, platform will have no choice but to encrypt on the fly, and write chiphertext blocks before all of the data is received, opening doors to blockwise adaptive attacks.

7 Implementation aspects

Software. In work to appear at CARDIS 2020, Deprez and coauthors make optimized software implementations of ForkAE (including ForkSkinny) on several platforms. Among other things, they show a decryption speed-up for implementations on low-end devices, and that SIMD hardware (e.g., x86 AVX or ARM Neon) can leverage multiple sources of parallelism in ForkSkinny.

Hardware. In our 2019 LWC Workshop paper, we already highlighted several interesting implementation strategies of the ForkSkinny primitive. In addition to

this, in his thesis Jowan Pittevils explored several speed-area trade-offs in the forkcipher. Among other things, he shows that forking can have very low impact on the implementation area (if necessary). More details can be found on the ForkAE website (`www.esat.kuleuven.be/cosic/forkae`).

8 Conclusion

The recent result summarized in this update document have revealed more of the true potential of the NIST LWC second round candidate ForkAE. Third party cryptanalysis, and our public challenge to break the novel ForkSkinny primitive have not revealed any unexpected weakness; on the contrary, they confirm the results of our own investigation. The new implementation results show the versatility and competitiveness of ForkAE, which can be efficiently implemented with various implementation trade-offs in SW and HW. The new security proofs show that the ForkAE family also gives much stronger guarantees than claimed previously, making it a candidate that covers "short-message", "general lightweight" and "defense in depth" use cases, and most importantly their intersection. We believe that this makes ForkAE a strong submission that can cater to the needs and constraints of real-world applications, while at the same time providing robust security.