Public Comments on the Draft SP 800-38G, Revision 1

(comment period closed April 15, 2019)

On February 28, 2019, NIST announced a period of public comment, ending April 15, 2019, on **Draft Special Publication 800-38G Revision 1, *Recommendation for Block Cipher Modes of Operation: Methods for Format-Preserving Encryption*** which updates the specifications of the FF1 and FF3 modes of the AES block cipher for format preserving encryption. The announcement was posted on the News and Events page at NIST's Computer Security Resource Center.

NIST received the following public comments:

David Gamey

Hi,

I have been reviewing the proposed update to 800-38G and have an observation about the proposed minimum domain size and how it might affect solutions in the credit card space.

In the use case where 16 digit credit cards are to be encrypted with FPE and the goal is to preserve the first 6, last 4, and Luhn. The middle six would be encrypted with the remaining 10 digits used to construct the tweak. As a result of the Luhn algorithm, which effectively rejects 90% of the domain, the domain size of the middle six digits will be reduced from 1M to 100K falling below the draft's minimum recommendation.

I have seen a number of (both proposed and implemented) solutions using format preserving techniques. These included not only FPE, but tokens, and random masking as well. There is a very strong tendency with these solutions to match Luhn. Almost all of the ones I have seen do so.

I can't say if reducing the domain strength from 1M to 100K is a good idea or not. I believe that NIST takes a generally conservative (safe) approach to such matters and there may be some room here for consideration. I'm also aware that research in this space is advancing quickly. Setting the domain strength to 1M will potentially create some awkward debates about the suitability of FPE for use with credit cards. Again, I am not taking a position either way but I believe that NIST should consider which course to take and provide comments on use cases with Luhn to ensure clarity.

Thanks,

David Gamey
CISSP, CISA, PCI: QSA(P2PE), PA-QSA(P2PE), 3DS Assessor, ISO/IEC: 27001 Lead Implementer
Senior Consultant & Researcher, Control Gap Inc.

Siddhartha Dutta

In regards to the request for public comments on the Draft Special Publication 800-38G (https://csrc.nist.gov/news/2019/nist-requests-comments-on-draft-sp-800-38g-rev-1), I wanted to share my thoughts as a consumer and implementer of the Format Preserving Encryption technology for several years across large financial institutions and global respected brands. I have reviewed the proposed revision to SP 800-38G and note a serious issue: it increases the smallest field that can be encrypted with the FF1 or FF3-1 modes from two decimal digits to six decimal digits. I understand the requirement to modify the FF1 and FF3-1 modes to account for the attacks that have been developed on short fields, but increasing the minimum field to six decimal digits is not the right way to do this IMHO.

I have always held a position around the value proposition of encrypting fields that are short such as State code, zip code, et al., and have challenged the need for it when they independently don't really reveal much about an individual if left in the clear, unless additional linked details are revealed as well. Choosing the right fields to encrypt and encrypting them using the right technology and model is more art than science. However, we cant avoid running into valid use cases where short fields are required to be encrypted, and by using format preserving encryption technology.

FPE is a very important and powerful technology. Many of the systems that process payments have legacy components that are either impossible or very expensive to modify to handle an encrypted value that does not have the same format at the corresponding plaintext. Thus, changing the SP 800-38G specification in a way that makes it incompatible with the most common use of the technology is a non-starter. Case in point, Primary Account Numbers (PANs, or credit card numbers) - This is one of the most widely used data element where FPE has been a game changer and have addressed enterprises' problems with implementing encryption on data, to cater to their data protection requirements, regulatory compliance needs, yet preserve business functions and database schema, and more importantly make encryption of credit cards, a transparent operation for any downstream applications within the business processes.

Although the ISO/IEC 7812-1 standard for PANs allow PANs to range from 12 to 19 digits, the vast majority of PANs currently used comprise 16 digits (and American Express has 15):
- The first six digits of a PAN specify the issuing bank for a card (called the "Issuer Identification Number" (IIN), or more colloquially, the "Bank Identification Number" (BIN).
- The next 9 digits are an account number at the issuing bank.
- The last digit is a checksum value.

The Payment Card Industry Data Security Standard (PCI DSS) that banks and merchants need to comply with allows the BIN and the last four digits of a PAN to remain unencrypted. This lends itself to the most common method to encrypt credit cards is to apply format preserving encryption partially on the data, whereby first 6 digits (BIN) and last 4 digits are left in the clear and the 6 digits in the middle are encrypted using FPE. This has tremendous value as this allows for routing applications to see the BIN in the clear during a payment transaction, or allow customer service applications to verify last 4 digits, or other applications to print portions of the

card on receipt or statements, etc., and all these without requiring these applications to decrypt the credit card number.

By 2022, the payments industry will move to 8-digit BINs while keeping 16-digit PANs. Longer PANs are not practical to implement because of the significant costs that would be involved in modifying existing payments networks to handle longer PANs. This leaves only the middle 4 digits of a PAN that can be encrypted. The proposed modification to SP 800-38G will not allow the modes that it specifies to be used to encrypt these middle four digits. This will eliminate the most important application of the technology. Because of this, NIST should consider other approaches that can increase the level of security that the modes described by SP 800-38G specify. Limiting the use of FPE to fields that comprise six or more digits will essentially make the technology useless.

Cryptographers may argue about the best way to increase the security of FPE when it is used to encrypt small fields. Some suggest that increasing the minimum field size that is allowed is the best solution. Others suggest that increasing the number of Feistel rounds is a reasonable alternative. While I am not a Cryptographer and do not have the specialized skills needed to assess the relative strengths and weaknesses of these two approaches, it is clear to me that increasing the minimum field size that can be FPE-encrypted to six digits will eliminate the most important use of the technology. Because of this, NIST should reconsider their approach to increasing the security of FPE when used to encrypt small fields and find a way that will keep the technology useful.

Looking at the initial NIST submissions for format-preserving encryption modes reveal an alternative to restricting the allowed size of fields that can be encrypted. In particular, the design to use more Feistel rounds to address the very issue that NIST is trying to address here: the reduced security that comes with smaller field sizes. That approach was to require more rounds for smaller field sizes. This likely makes the encryption algorithm more complex, but it allows the technology to be used in its most important use case. Performance of FPE operations has never really been a topic of concern, especially when we are dealing with the order of 10s of microseconds, and hence adding additional rounds of Feistel or increased complexity of the encryption operation, ONLY during the cases of shorter fields, even if it doubles or triples the overhead, that should not be a concern for the industry at large.

To summarize, increasing the minimum field size that can be encrypted with the FF1 and FF3-1 algorithms to six digits should be avoided. Doing so would make the technology useless for its most important use case, encrypting credit card numbers. NIST should explore other approaches that address the small-domain weakness, instead of limiting use of the technology such that it becomes useless for most existing users.

Regards,
Sid Dutta.

Kevin Hamilton

Comments for consideration regarding SP800-38G Rev 1.

The proposed revision to SP 800-38G could have a negative impact to data security. It increases the smallest field that can be encrypted with the FF1 or FF3-1 modes from two decimal digits to six decimal digits. I understand the requirement to modify the FF1 and FF3-1 modes to account for the attacks that have been developed on short fields. However, increasing the minimum field to six decimal digits may not be the best way to do this. It could render sound field level encryption practices out of NIST compliance.

Most format preserving field level encryption for a CCN leaves the first six digits (BIN)and the last four digits in the clear. The BIN is used in the clear for routing purposes. The last four is used to print statements and for customer service representatives to authenticate users. This allows the data to remain encrypted even while being used for most business processes – a very good thing for data security.

The industry will be moving to longer BINs in the near future. With the first eight digits in the clear for extended BINs, that would only leave the four middle digits encrypted unless CCNs are extended beyond that traditional 15-16 digits. This does not seem likely in the near term due to the amount of impact it would have on the card associations, processors, banks and merchants.

Unfortunately, the proposed revision to SP 800-38G would not allow the use of the FF1 or FF3-1 modes to encrypt fields of only four decimal digits. Providing security for credit card numbers is one of the most important use cases for the modes specified by SP 800-38G. Eliminating this use case is something that could negatively impact data security for PCI data.

There are alternatives to restricting the field size that can be encrypted with FPE in order to address the reduced security that comes with a smaller field size. For example, If more rounds of encryption were required for smaller field sizes, that would make the encryption more complex to solve. That approach would allow for the continued use of FPE, and preserve the additional data security associated with keeping the data encrypted even while it is in use.

Thank you for your consideration.

Sincerely,

Kevin Hamilton

Kevin Hamilton
Vice President, Global Cyber Security
First Data

The attacks on FF1 and FF3 exploit the following issues:

- The aggressive optimization of tweaks in FF3 ([DV17,HMT19]). The current fix in FF3.1 is to require some particular 8 bits of the tweak to be zero. This is enough to thwart the attacks above.
- The round count is not enough for tiny domains [BHT16,HTT17,DV18]. The current fix in both FF1 and FF3.1 is to require that the domain size must be at least 10^6. This would make the attack in [DV18] more expensive than a brute-force attack on the key. On the other hand, between [BHT16] and [HTT17], the latter requires weaker assumption on the known-message distribution, and better *amortized* cost if one aims to recover multiple target messages. However, here the "right" metric is to look at the cost of recovering a single target message, and make sure that it's prohibitive. In that metric, the Left-Half attack in [BHT16] is the best. Below, I give an estimate of data complexity of that attack on FF1 and FF3.1 for two domains: 6-decimal-digit and 7-decimal-digit. The timing cost is around the same as the data cost.

Scheme	Data (6 digits)	Data (7 digits)
FF1	2^{77}	2^{100}
FF3.1	2^{57}	2^{54}

Clearly, the cost for FF1 is prohibitive: the adversary can't hope to gather much data in practice. For FF3.1, the cost for 7-digit domain is *smaller* than that of the 6-digit domain, because FF3.1 has an inferior way of partitioning odd domain size, compared to FF1. I leave it to NIST to decide whether the cost of FF3.1 is prohibitive in practice. Note that the attack in [BHT16] still works even if one re-keys often.

Jay Irwin

Ladies and gentlemen:

The following are my comments as an individual and are not made on behalf of Teradata Corporation or any of its subsidiaries or entities

I do not see that sufficient justification has been given in the current proposed revision for FF3 to continue, as is (noting the April 2017 guidance) or as FF3-1 as in the current proposed revision.

I support the increase in domain size requirements for the reasons given.

Jay Irwin, JD
Director, Teradata Center for Enterprise Security

Eric Lengvenis

To whom it may concern,

I would like to offer a statement of support for the effort to standardize the three format-preserving modes of operation in the current draft of SP 800-38G and to make a few minor comments. As a large financial institution we have a preference for technology in conformance with standards put out by NIST and ASC X9. These help guarantee a level of confidence in the implementation of encryption technology. To this end, we have been working with X9 to standardize FPE in X9.124, but this would be even more valuable as NIST standards if the validation of the modes is incorporated into the FIPS 140 validation program. If this comes to be, our HSM vendors could incorporate FPE into FIPS-compliant appliances which is very desirable. Already FPE is widely-used but not defined in a standard which creates tension between using the technology that allows us to encrypt in legacy applications which cannot be overhauled to allow for the format changes required by other approaches and our preference for standardized technology. For these reasons, I support this effort.

On to the comments. There are two source errors in the document; one on page 15 and the other on page 17. Both appear to be referencing the Feistel diagram, but the link is incorrect. The other, is a question -- why is not the full BPS approved, but the specific subset is? I think it would merit a statement as to why, given that the proposed mode defines it.

Thank you,

Eric Lengvenis
Information Security Architect
Vice President
Enterprise Information Security Architecture (EISA)

Andrew Price

We fully support moving forward with the publication of SP 800-38G. The format-preserving encryption technologies that it specifies are an important tool for protecting sensitive information in complex IT environments, and their availability can make the difference between sensitive data being encrypted and sensitive data not being encrypted. It's hard to get an accurate estimate of the economic losses caused by data breaches, but it's certain that these losses can be greatly reduced by the more widespread use of encryption. And because the technologies defined by SP 800-38G make this practical when it would not be practical with existing encryption approaches, making these technologies acceptable for broad use is definitely a step in the right direction.

Regards,

Andrew Price

Director, Product Management
XYPRO Technology Corporation
HP NonStop Server Security
and Encryption Solutions