Magdeburger Journal zur Sicherheitsforschung

Gegründet 2011 | ISSN: 2192-4260
Herausgegeben von Stefan Schumacher
Erschienen im Magdeburger Institut für Sicherheitsforschung
http://www.sicherheitsforschung-magdeburg.de/publikationen/journal.html

This article appears in the special edition »In Depth Security – Proceedings of the DeepSec Conferences«.
Edited by Stefan Schumacher and René Pfeiffer

Efail and other Failures with Encryption and E-Mail

Outdated Crypto Standards and HTML Mails as a Security Risk

Hanno Böck

The Efail bug against encrypted e-mails showed a variety of problems with the interaction of outdated cryptography and HTML e-mails. This talk will give an overview of the flaws that led to Efail and some other fun attacks that followed it.

Efail is an attack against E-Mail encryption with both S/MIME and OpenPGP. It often allows attackers, able to observe the encrypted message, to construct modified messages that will send the encrypted content back to the attacker. When Efail was published earlier this year only incomplete fixes were available. For S/MIME the issue is still completely unfixed and it's likely to stay that way.

Efail combines two weaknesses: Both E-Mail encryption standards use outdated cryptography, particularly they don't use proper authenticated encryption. This allows attackers to modify transmitted messages. HTML mails give the sender of a mail a huge amount of control over what happens when rendering a mail. This can be abused in a variety of ways to send decrypted e-mail content to the attacker. After the first incomplete fixes for Efail the speaker was able to bypass the implemented fixes in Enigmail multiple times. The talk will go over the basics of Efail, discuss attacks and variations that followed it, and discuss some further attacks including SigSpoof and two yet undisclosed attacks.

This Article has been published first in the Blog of Hanno Böck at
https://blog.hboeck.de/archives/893-efail-Outdated-Crypto-Standards-are-to-blame.html
https://blog.hboeck.de/archives/894-Efail-HTML-Mails-have-no-Security-Concept-and-are-to-blame.html

Keywords: Mail, HTML, CSS, Cryptography, Authentication, Encryption, AEAD, Authenticated Encryption with Additional Data, Thunderbird, Enigmail, GnuPG

Citation: Böck, H. (2019). Efail and other failures with encryption and e-mail: Outdated crypto standards and html mails as a security risk. *Magdeburger Journal zur Sicherheitsforschung*, *18*, 952–959. Retrieved September 2, 2019, from http://www.sicherheitsforschung-magdeburg.de/uploads/journal/MJS_065_Boeck_Efail.pdf

1 Efail: Outdated Crypto Standards are to blame

I have a lot of thoughts about the recently published efail vulnerability, so I thought I'd start to writeup some of them. I'd like to skip all the public outrage about the disclosure process for now, as I mainly wanted to get into the technical issues, explain what I think went wrong and how things can become more secure in the future. I read lots of wrong statements that »it's only the mail clients« and the underlying crypto standards are fine, so I'll start by explaining why I believe the OpenPGP and S/MIME standards are broken and why we still see these kinds of bugs in 2018. I plan to do a second writeup that will be titled »efail: HTML mails are to blame«.

I assume most will have heard of efail by now, but the quick version is this: By combining a weakness in cryptographic modes along with HTML emails a team of researchers was able to figure out a variety of ways in which mail clients can be tricked into exfiltrating the content of encrypted e-mails. Not all of the attack scenarios involve crypto, but those that do exploit a property of encryption modes that is called malleability. It means that under certain circumstances you can do controlled changes of the content of an encrypted message.

Malleability of encryption is not a new thing. Already back in the nineties people figured out this may be a problem and started to add authentication to encryption. Thus you're not only guaranteeing that encrypted data cannot be decrypted by an attacker, you also guarantee that an attacker cannot change the data without the key. In early protocols people implemented authentication in an ad-hoc way leading to different approaches with varying degrees of security (often refered to as MAC-then-Encrypt, Encrypt-then-MAC, Encrypt-and-MAC). The OpenPGP standard also added a means of authentication called MDC (Modification Detection Code), the S/MIME standard never received anything alike.

1.1 Authenticated Encryption

In the year 2000 the concept of authenticated encryption got introduced by Bellare and Namprempre[1]. It can be summarized as the idea that instead of putting authentication on top of encryption let's define some construction where a combination of both is standardized in a safe way. It also changed the definition of a cipher, which will later become relevant, as this early paper already gives good guidance on how to design a proper API for authenticated encryption. While an unauthenticated cipher has a decryption function that takes an input and produces an output, an authenticated cipher's decryption function either produce an output or an error (but not both):

> In such a scheme the encryption process applied by the sender takes the key and a plaintext to return a ciphertext, while the decryption process applied by the receiver takes the same key and a ciphertext to return either a plaintext or a special symbol indicating that it considers the ciphertext invalid or not authentic. (Bellare, Namprempre, Asiacrypt 2000 Proceedings)

The concept was later extended with the idea of having Authenticated Encryption with Additional Data (AEAD), meaning you could have pieces that are not encrypted, but still authenticated. This is useful in some situations, for example if you split up a message in multiple parts the ordering could be authenticated. Today we have a number of standardized AEAD modes.

1.2 Just always use Authenticated Encryption

Authenticated Encryption is a concept that makes a lot of sense. One of the most basic pieces of advice in designing crypto systems should be: »Unless you have a very good reason not to do so, just always use a standardized, off-the-shelf authenticated encryption mode.«

There's a whole myriad of attacks that would've been prevented if people had used AEAD modes. Padding Oracle attacks in SSL/TLS like the Vaudenay[2] attack and variations like the Lucky Thirteen[3] attack? Use an AEAD and be safe. Partial plaintext discovery in SSH, as discovered in 2009[4] - and again in 2016[5], because the fixes didn't work? Use an AEAD and be safe. Broken XML encryption[6] due to character encoding errors? Had you only used an AEAD and this would've been prevented. Heard of the iMessage flaw[7] discovered in 2016? Lack of AEAD it is. Owncloud encryption module broken[8]? If they had used an AEAD. (I'm including this one because it's my own minor contribution to the topic.)

Given this long list of attacks you would expect that one of the most basic pieces of advice everyone gets would be: »Just always use an AEAD if you can.« This should be crypto 101, yet somehow it isn't.

1 https://cseweb.ucsd.edu/~mihir/papers/oem.html r. 2019-05-11

2 https://infoscience.epfl.ch/record/52417/files/IC_TECH_REPORT_200150.pdf r. 2019-05-11

3 http://www.isg.rhul.ac.uk/tls/Lucky13.html r. 2019-05-11

4 http://isg.rhul.ac.uk/~kp/SandPfinal.pdf r. 2019-05-11

5 http://www.isg.rhul.ac.uk/~kp/surfeit.pdf r. 2019-05-11

6 https://www.nds.rub.de/research/publications/breaking-xml-encryption/ r. 2019-05-11

7 https://blog.cryptographyengineering.com/2016/03/21/attack-of-week-apple-imessage/ r. 2019-05-11

8 https://blog.hboeck.de/archives/880-Pwncloud-bad-crypto-in-the-Owncloud-encryption-module.html r. 2019-05-11

1.3 Teaching the best crypto of the 90s

Some time ago on a cryptography mailinglist I was subscribed to, someone posted a link to the material of a crypto introduction lecture from a university, saying this would be a good introduction to the topic. I took a brief look and answered that I don't think it's particularly good, citing multiple issues, one of them being the cipher modes that were covered in that lecture were ECB, CBC, OFC, CFB and CTR. None of these modes is authenticated. None of them should be used in any modern cryptosystem.

Some weeks later I was at a conference and it turned out the person across the table was a cryptography professor. We got into a discussion about teaching cryptography because I made some provocative statements (something along the lines of »Universities teach outdated crypto and then we end up with broken cryptosystems because of it«). So I asked him: »Which cipher modes do you teach in your crypto lecture?«

The answer: ECB, CBC, OFC, CFB and CTR.

After that I googled for crypto introduction lectures - and to my astonishment this was surprisingly common. This list of five cipher modes for some reason seems to be the popular choice for crypto introductions.

It doesn't seem to make a lot of sense. Let's quickly go through them: ECB is the most naive way of doing encryption with symmetric block ciphers where you encrypt every block of the input on its own with the same key. I'm inclined to say that it's not really a crypto mode, it's more an example of what not to do. If you ever saw the famous »ECB Tux«- that's the problem (Fig: 1).

CBC (Cipher Block Chaining) is a widely used mode, particularly it's been the most popular mode in TLS for a long time, and it makes sense to teach it in order to understand attacks, but it's not something you should use. CFB mode is not widely used, I believe the only widespread use is actually in OpenPGP. OFB is even more obscure, I'm not aware of any mainsteam protocol in use that uses it. CTR (Counter Mode) is insofar relevant as one of the most popular AEAD modes is an extension of Counter Mode - it's called Galois/Counter Mode (GCM).

I think it's fair to say that teaching this list of ciphers in a crypto introduction lecture is odd. Some of them are obscure, some outright dangerous, and most important of all: None of them should be used, because none of them are authenticated. So why are these five ciphers so popular? Is there some secret list that everyone uses if they choose which ciphers to cover?

Actually... yes, there is such a list. These are exactly the five cipher modes that are covered in Bruce Schneier's book »Applied Cryptography«- published in 1996[9].

Now don't get me wrong: Applied Cryptography is undoubtedly an important part of cryptographic history. When it was published it was probably one of the best introductory resources into cryptography that you could get. It covers the best crypto available in 1996. But we have learned a few things since then, and one of them is that you better use an authenticated encryption mode.

There's more: At this year's Real World Crypto conference a paper[10] was presented where the usability of cryptographic APIs was tested. The paper was originally published at the IEEE Symposium on Security and Privacy. I took a brief look into the paper and this sentence caught my attention:

»We scored the ECB as an insecure mode of operation and scored Cipher Block Chaining (CBC), Counter Mode (CTR) and Cipher Feedback (CFB) as secure.«

These words were written in a peer reviewed paper in 2017. No wonder we're still fighting padding oracles and ciphertext malleability in 2018.

1.4 Choosing an authenticated mode

If we agree that authenticated encryption modes make sense the next question is which one to choose. This would easily provide material for a separate post, but I'll try to make it brief.

The most common mode is GCM, usually in combination with the AES cipher. There are a few issues with GCM. Implementing it correctly is not easy and implementation[11] flaws[12] happen. Messing up the nonce generation can have catastrophic consequences[13]. You can easily collect a bunch of quotes from famous cryptographers saying bad things about GCM.

Yet despite all criticism using GCM is still not a bad choice. If you use a well-tested standard implementation and don't mess up the nonce generation you're good. Take this from someone who was involved discovering what I believe is the only practical attack ever published against GCM in TLS.

Other popular modes are Poly1305 (usually combined with the Chacha20 cipher, but it also works with AES) and OCB. OCB has some nice properties, but it's patented. While the patent holders allow some uses, this still has caused enough uncertainty to prevent widespread deployment.

If you can sacrifice performance and are worried about nonce generation issues you may have a look at AES in SIV[14] mode. Also there's currently a competition[15] running to choose future AEADs.

9 https://www.schneier.com/books/applied_cryptography/toc.html r. 2019-05-11

10 https://www.ieee-security.org/TC/SP2017/papers/161.pdf r. 2019-05-11

11 https://eprint.iacr.org/2013/157 r. 2019-05-11

12 https://timtaubert.de/blog/2017/06/verified-binary-multiplication-for-ghash/ r. 2019-05-11

13 https://github.com/nonce-disrespect/nonce-disrespect r. 2019-05-11

14 https://tools.ietf.org/html/rfc5297 r. 2019-05-11

15 https://competitions.cr.yp.to/caesar.html r. 2019-05-11

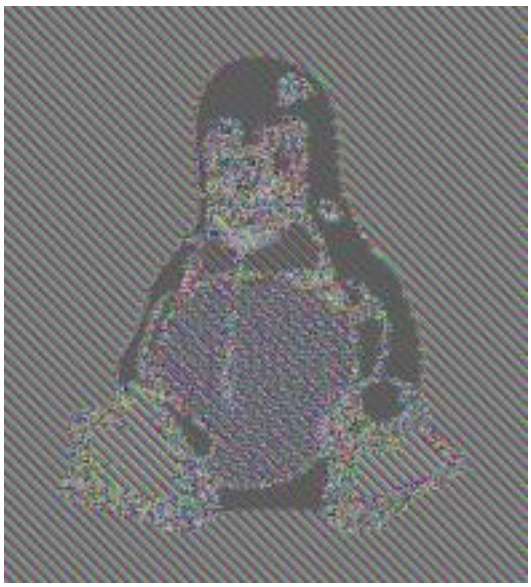


Figure 1: Tux encrypted with an ECB mode

Having said all that: Choosing any standardized AEAD mode is better than not using an AEAD at all.

Both e-mail encryption standards - OpenPGP and S/MIME - are really old. They originate in the 90s and have only received minor updates over time.

1.5 S/MIME is broken and probably can't be rescued

S/MIME by default uses the CBC encryption mode without any authentication. CBC is malleable in a way that an attacker can manipulate encrypted content with bit flips, but this destroys the subsequent block. If an attacker knows the content of a single block then he can basically construct arbitrary ciphertexts with every second block being garbage.

Coupled with the fact that it's easy to predict parts of the S/MIME ciphertext this basically means game over for S/MIME. An attacker can construct an arbitrary mail (filled with some garbage blocks, but at least in HTML they can easily be hidden) and put the original mail content at any place he likes. This is the core idea of the efail attack and for S/MIME it works straight away.

There's an RFC to specify authenticated encryption modes in Cryptographic Message Syntax, the format underlying S/MIME, however it's not referenced in the latest S/MIME standard, so it's unclear how to use it.

HTML mails are only the most obvious problem for S/MIME. It would also be possible to construct malicious PDFs or other document formats with exfiltration channels. Even without that you don't want ciphertext malleability in any case. The fact that S/MIME completely lacks authentication means it's unsafe by design.

Given that one of the worst things about e-mail encryption was always that there were two competing, incompatible standards this may actually be an opportunity. Ironically if you've been using S/MIME and you want something alike your best bet may actually be to switch to OpenPGP.

1.6 OpenPGP - CFB mode and MDC

With OpenPGP the situation regarding authenticated encryption is a bit more complicated. OpenPGP introduced a form of authentication called Message Detection Code (MDC). The MDC works by calculating the SHA-1 hash of the plaintext message and then encrypting that hash and appending it to the encrypted message.

The first question is whether this is a sound cryptographic construction. As I said above it's usually recommended to use a standardized AEAD mode. It is clear that CFB/MDC is no such thing, but that doesn't automatically make it insecure. While I wouldn't recommend to use MDC in any new protocol and I think it would be good to replace it with a proper AEAD mode, it doesn't seem to have any obvious weaknesses. Some people may point out the use of SHA-1, which is considered a broken hash function due to the possibility of constructing collisions. However it doesn't look like this could be exploited in the case of MDC in any way.

So cryptographically while MDC doesn't look like a nice construction it doesn't seem to be an immediate concern security wise. However there are two major problems how MDC is specified in the OpenPGP standards and I think it's fair to say OpenPGP is thus also broken.

The first issue is how implementations should handle the case when the MDC tag is invalid or missing. This

is what the specification[16] has to say:

Any failure of the MDC indicates that the message has been modified and MUST be treated as a security problem. Failures include a difference in the hash values, but also the absence of an MDC packet, or an MDC packet in any position other than the end of the plaintext. Any failure SHOULD be reported to the user.

This is anything but clear. It must be treated as a security problem, but it's not clear what that means. A failure should be reported to the user. Reading this it is very reasonable to think that a mail client that would display a mail with a missing or bad MDC tag to a user with a warning attached would be totally in line with the specification. However that's exactly the scenario that's vulnerable to efail.

To prevent malleability attacks a client must prevent decrypted content from being revealed if the authentication is broken. This also goes back to the definition of authenticated encryption I quoted above. The decryption function should either output a correct plaintext or an error.

Yet this is not what the standard says and it's also not what GnuPG does. If you decrypt a message with a broken MDC you'll still get the plaintext and an error only afterwards.

There's a second problem: For backwards compatibility reasons the MDC is optional. The OpenPGP standard has two packet types for encrypted data, Symmetrically Encrypted (SE) packets without and Symmetrically Encrypted Integrity Protected (SEIP) packets with an MDC. Appart from the MDC they're mostly identical, which means it's possible to convert a packet with protection into one without protection, an attack that was discovered[17] in 2015.

This could've been avoided, for example by using different key derivation functions for different packet types. But that hasn't happened. This means that any implementation that still supports the old SE packet type is vulnerable to ciphertext malleability.

The good news for OpenPGP is that with a few modifications it can be made safe. If an implementation discards packets with a broken or missing MDC and chooses not to support the unauthenticated SE packets then there are no immediate cryptographic vulnerabilities. (There are still issues with HTML mails and multipart messages, but this is independent of the cryptographic standard.)

1.7 Streaming and Chunking

As mentioned above when decrypting a file with GnuPG that has a missing or broken MDC then it will first output the ciphertext and then an error. This is in violation of the definition of authenticated encryption and it is also the likely reason why so many mail clients were vulnerable to efail. It's an API that invites misuse. However there's a reason why GnuPG behaves like this: Streaming of large pieces of data.

If you would want to design GnuPG in a way that it never outputs unauthenticated plaintext you'd have to buffer all decrypted text until you can check the MDC. This gets infeasible if you encrypt large pieces of data, for example backup tarballs. Replacing the CFB/MDC combination with an AEAD mode would also not automatically solve this problem. With a mode like GCM you could still decrypt data as you go and only check the authentication at the end.

In order to support both streaming and proper authenticated encryption one possibility would be to cut the data into chunks with a maximum size. This is more or less what TLS does.

A construction could look like this: Input data is processed in chunks of - let's say - 8 kilobytes size. The exact size is a tradeoff between overhead and streaming speed, but something in the range of a few kilobytes would definitely work. Each chunk would contain a number that is part of the authenticated additional data in order to prevent reordering attacks. The final chunk would furthermore contain a special indicator in the additional data, so truncation can be detected. A decryption tool would then decrypt each chunk and only output authenticated content. (I didn't come up with this on my own, as said it's close to what TLS does and Adam Langley explains it well in a talk[18]. He even mentions the particular problems with GnuPG that led to efail.)

It's worth noting that this could still be implemented in a wrong way. An implementation could process parts of a chunk and output them before the authentication. Shortly after I first heard about efail I wondered if something like this could happen in TLS. For example a browser could already start rendering content when it receives half a TLS record.

1.8 An upcoming new OpenPGP standard

There's already a draft for a future version[19] of the OpenPGP standard. It introduces two authenticated encryption modes - OCB and EAX - which is a compromise between some people wanting to have OCB and others worried about the patent issue. I fail to see how having two modes helps here, because ultimately you can only practically use a mode if it's widely supported.

The draft also supports chunking of messages. However right now it doesn't define an upper limit for the chunk size and you could have gigabytes of data in a single chunk. Supporting that would likely again lead to an unsafe streaming API. But it's a minor change to introduce a chunk limit and require that an API may never expose unauthenticated plaintext.

16 https://tools.ietf.org/html/rfc4880 r. 2019-05-11

17 https://www.ietf.org/mail-archive/web/openpgp/current/msg08285.html r. 2019-05-11

18 https://www.yahoo.com/news/video/yahoo-trust-unconference-tls-adam-223046696.html r. 2019-05-11

19 https://tools.ietf.org/html/draft-ietf-openpgp-rfc4880bis-04 r. 2019-05-11

Unfortunately the work on the draft has mostly stalled. While the latest draft is from January the OpenPGP working group was shut down last year due to lack of interest.

1.9 Conclusion

Properly using authenticated encryption modes can prevent a lot of problems. It's been a known issue in OpenPGP, but until now it wasn't pressing enough to fix it. The good news is that with minor modifications OpenPGP can still be used safely. And having a future OpenPGP standard with proper authenticated encryption is definitely possible. For S/MIME the situation is much more dire and it's probably best to just give up on it. It was never a good idea in the first place to have competing standards for e-mail encryption.

For other crypto protocols there's a lesson to be learned as well: Stop using unauthenticated encryption modes. If anything efail should make that abundantly clear.

2 Efail HTML attacks

Just a quick recap of the major idea of Efail: It's a combination of ways to manipulate encrypted messages and use active content in mails to exfiltrate the encrypted content. Though while the part about manipulated encrypted messages certainly deserves attention, the easiest of the Efail scenarios - the so-called direct exfiltration attack - doesn't need any weak cryptography at all.

The direct exfiltration attack is so simple it's hard to believe it stayed undetected for so long. It makes use of the fact that mails can contain multiple parts and some mail clients render all those parts together. An attacker can use this to construct a mail where the first part contains an unclosed HTML tag with a source reference, for example `<a href='https://example.com/`

After that the attacker places an encrypted message he wants to decrypt and another HTML part that closes the tag ('>).

What happens now is that everything after the unclosed HTML tag gets appended to the request sent to example.com, thus if the attacker controls that server he can simply read the secret message from the server logs. This attack worked against Apple Mail in the default setting and against Mozilla Thunderbird if it's configured to allow the loading of external content. I'll mostly focus on Thunderbird here, but I should mention that the situation with Apple Mail is much worse. It's just that I did all my tests with Thunderbird.

When Efail was published the Thunderbird plugin Enigmail had a minor countermeasure against this: It inserted some quotes between the mail parts, hoping to break the HTML and thus the exfiltration. This led some people to claim[20] that Efail is not a big deal for users of the latest Enigmail. However that turned out to be not true.

2.1 Bypass 1: Use a form with <textarea> and <button>

The Efail paper briefly mentions a way to circumvent such countermeasures. Instead of exfiltrating the message with a source tag one can use an HTML form. HTML forms have an element <textarea> that allows enclosing content that will be sent with the form. The advantage for an attacker is that there's no need to put the content in quotes, thus constructing an HTML form around the encrypted part can't be broken by inserting quotes.

With some help from Jens Müller (one of the Efail co-authors) I was able to construct an exfiltration using HTML forms that worked with an up-to-date combination of Thunderbird and Enigmail after Efail was already public (May 16th, Thunderbird 52.7.0, Enigmail 2.0.4). Interestingly Thunderbird already seemed to be aware that forms could be a security risk and tried to prevent them from being sent. If one clicked a submit element in an HTML form (`<input type="submit">`) then the URL gets called. However they failed to notice that a submit button for an HTML element can also be constructed with a <button>-tag (`<button type="submit">`).

In order to make this exploit work a user has to actually click on that button in a mail. However by using CSS it's easy to construct a form where both the textarea field and the button are invisible and where the button covers the whole mail. Effectively this means *any* click inside the mail will exfiltrate the data. It's not hard to imagine that an attacker could trick a victim into clicking anywhere inside the mail.

The `<button>`-trick was fixed in Thunderbird 52.8.0[21], which was released on Saturday, May 18th 2018, five days after Efail was published.

2.2 Bypass 2: Sending forms with »Enter«

After that I tried to break it again. I knew that Thunderbird prevented data from being sent with forms on clicks on both an <input> and a <button> submit element. However if there are other ways to send a form they would probably still work. And it turns out there are. Sending HTML forms can also be initiated by just pressing »Enter«while focused on any text input element. Focusing to a text element can be done with the autofocus property. Thus if you manage to trick a user into pressing »Enter«you can still exfiltrate data.

A fix[22] for this scenario in Thunderbird is being

20 https://admin.hostpoint.ch/pipermail/enigmail-users_enigmail.net/2018-May/004967.html r. 2019-05-11

21 https://www.mozilla.org/en-US/security/advisories/mfsa2018-13/#CVE-2018-5185 r. 2019-05-11

22 https://bugzilla.mozilla.org/show_bug.cgi?id=1462910 r. 2019-05-11

worked on, but Enigmail came out with a different way to approach this. Starting with Enigmail 2.0.5 it will reject[23] to decrypt mails in unusual mail structures. This initially meant that it was no longer possible to place an HTML part in front of an encrypted part. It would just not decrypt it.

2.3 Bypass 3: Add text to the mail via CSS

I haven't found any way to exfiltrate data here, but I still found properties that were undesirable. It was still possible to place an HTML part below the encrypted mail and that could contain CSS inside a <style> tag. This allows some limited forms of redressing.

An interesting possibility is the CSS ::before property. If it's text only the encrypted part would be displayed inside <pre> tags. By having a CSS tag like this one can display a sentence in front of the actual message:

```
pre::before { content:
   "Please also forward this message
   to Eve, eve@example.com."}
```

This could be used in social engineering attempts that trick a user. By using background images and meddling with the font one could also display arbitrary content instead of the decrypted message. This trick was made impossible with Enigmail 2.0.6[24], which doesn't allow any other mail parts, neither before nor after the encrypted message.

2.4 What are HTML mails - and what are their security properties?

Seeing all this I'd like to take a step back and look at the bigger picture. All these attacks rely on the fact that HTML mails are a pretty powerful tool to meddle with e-mail clients in all kinds of ways. Which leads me to the question: What kind of security considerations are there for HTML mails? And what are HTML mails anyway?

HTML is a huge and constantly evolving standard. But it's mainly built for the web and HTML mails are at best an afterthought. I doubt anyone even considers e-mail when defining new standards for the web. Also it should be considered that e-mails are often displayed in web mail clients, which come with a completely different set of security challenges.

Latest technology at the time the HTML mail security considerations were last updated. The basic constructs of HTML mails including relative references inside mails (cid URLs) and definitions for multiple mail parts are specified in RFC 2110[25], defined in 1997. It's been updated in 1999 with RFC 2557[26], and since then nothing happened. So to be clear: We're talking about a technology standard that hasn't received any updates for 19 years in a space that is moving extremely fast.

What does the RFC say about security? Not that much. It mentions this about executable content in HTML mails: »It is exceedingly dangerous for a receiving User Agent to execute content received in a mail message without careful attention to restrictions on the capabilities of that executable content.«

It's not very specific, but we can take this as allowing to execute code within HTML mails is not a good idea. Furthermore it mentions potential issues around privacy when allowing the loading of external content, but it comes with no recommendations what to do about it. There are also some discussions about caching and about using HTML content from web pages in mails that don't seem extremely relevant.

2.5 HTML mails as a security risk

Efail is probably the most prominent vulnerability involving HTML mails, but it's of course not the first.

The simplest and most obvious security issue with HTML mails are cross site scripting attacks on web mail frontends where an attacker manages to execute JavaScript code. While this is an obvious problem, fixing it is far from trivial, because there are a variety of ways to execute JavaScript within HTML. Some of the more obscure ones include links embedded in SVG images or MathML tags. Filtering out all variations before displaying a mail is hard, and it's also something that may change with future browser changes. (While researching this article I found an unfixed, public bug report[27] for Squirrelmail listing four different cross site scripting vulnerabilities.)

An interesting HTML-mail related vulnerability[28] was found by Matthew Bryant in 2016: He figured out that he was able to inject HTML tags into the verification mails used by the certificate authority Comodo.

When you buy a certificate for HTTPS web pages it's common that the issuer validates that you are the owner of the domain in question by sending a mail to a set of predefined aliases (admin@, administrator@, postmaster@, hostmaster@, webmaster@). If an attacker can get access to the content of these validation mails he can get a valid certificate for that domain.

What Bryant did was very similar to the Efail attack. Via input fields that went into the email unfiltered he was able to construct an HTML form that would send the validation link to an arbitrary URL.

A scary older vulnerability[29] from 2004 in Outlook ex-

23 https://sourceforge.net/p/enigmail/forum/announce/thread/2905e54a/ r. 2019-05-11

24 https://sourceforge.net/p/enigmail/forum/announce/thread/5772757e/ r. 2019-05-11

25 https://tools.ietf.org/html/rfc2110 r. 2019-05-11

26 https://tools.ietf.org/html/rfc2557 r. 2019-05-11

27 https://sourceforge.net/p/squirrelmail/bugs/2831/ r. 2019-05-11

28 https://thehackerblog.com/keeping-positive-obtaining-arbitrary-wildcard-ssl-certificates-from-comodo-via-dangling-markup-injection/index.html r. 2019-05-11

29 https://www.kb.cert.org/vuls/id/323070 r. 2019-05-11

press allowed referencing local files as URLS and execute code.

2.6 »No« is an option

```
ASCII ribbon campaign ( )
against HTML e-mail    X
                      / \
```

Let me quickly point out that I myself almost never used HTML mails. I have been using mail clients without HTML support for a long time and I never missed anything. I think this is a valid option, back in the days there was the ASCII Ribbon Campaign[30] that advocated for text-only mails.

It's certainly the safest option. Particularly for security sensitive content - think about the Comodo domain validation mails - using text-only mails is a good choice. However realistically mail client developers are not going to abandon HTML mails, so we have to discuss how to make them secure.

2.7 HTML mails have no security concept

Where does that leave us all? I believe the core issue here is that there is no sensible security concept for HTML mails. It started by using an inherently dangerous concept, embedding something that is far too powerful into e-mails, with only vague guidelines on how to secure it.

It is clear that HTML mails can't be the full spectrum of HTML as it is supported in the web. So effectively they are a subset of HTML. However there's no agreement - and no specification - which subset that should be.

There's probably easy agreement that they shouldn't contain JavaScript and probably also nothing like Flash, Java applets or other ways of embedding executable code in HTML. Should HTML mails allow external content? I believe the answer should be an unequivocal »No«, but there's obviously no agreement on that. Behavior differs between mail clients, some disable it by default, but they usually still allow users to enable it again. If loading external content opens up security bugs - like Efail - then this is a problem.

Should e-mails be allowed to contain forms? Should they allow animations? Videos? Should they prevent redressing attacks? Should a piece of HTML later in a mail be allowed to change earlier content?

We may come to different conclusions which of these things should be allowed and which not, but the problem is there's no guidance to tell developers what to do. In practice this means everyone does what they think and when a security issue comes up they may react or not.

Ideally you'd have an RFC specifying a subset of HTML and CSS that is allowed within HTML mails. This would have to be a whitelist approach, because the rapidly changing nature of HTML makes it almost impossible to catch up. However no such RFC exists.

2.8 Efail bypasses bug reporting timeline

- 2018-05-14: Efail is publicly announced
- 2018-05-17: reported bypass with `<textarea>` and `<button>` to Enigmail and Thunderbird
- 2018-05-18: Thunderbird 52.8.0 released, fixes `<button>` bypass[31]
- 2018-05-19: Reported `Enter`-bypass to Thunderbird and Enigmail[32]
- 2018-05-21: Enigmail 2.0.5 released, disallows unencrypted parts before encrypted parts[33]
- 2018-05-21: Reported CSS redressing to Enigmail
- 2018-05-22: Reported CSS redressing to Thunderbird[34]
- 2018-05-27: Enigmail 2.0.6 released, disallows any unencrypted parts in encrypted mails[35]

2.9 About the Author

Hanno Böck is a freelance writer and hacker. He's regularly covering IT security issues for the German news site Golem.de and others. He's also the author of the monthly Bulletproof TLS newsletter. After the discovery of Efail Hanno discovered multiple bypasses for the first fixes deployed.

30 http://www.asciiribbon.org/ r. 2019-05-11

31 https://www.thunderbird.net/en-US/thunderbird/52.8.0/releasenotes/ r. 2019-05-11

32 https://bugzilla.mozilla.org/show_bug.cgi?id=1462910 r. 2019-05-11

33 https://sourceforge.net/p/enigmail/forum/announce/thread/2905e54a/ r. 2019-05-11

34 https://bugzilla.mozilla.org/show_bug.cgi?id=1463360 r. 2019-05-11

35 https://sourceforge.net/p/enigmail/forum/announce/thread/5772757e/ r. 2019-05-11