Security Now! #670 - 07-03-18

Wi-Fi Protected Access v3

This week on Security Now!

This week we discuss the interesting case of a VirusTotal upload... or was it?, newly discovered problems with our 4G LTE... and even what follows, another new EFF encryption initiative, troubles with Spectre and Meltdown in some browsers, the evolution of UPnP-enabled attacks, an unpatched Wordpress vulnerability that doesn't appear to be worrying the Wordpress devs... and an early look at next year's forthcoming WPA3 standard... which appears to fix everything!

Our Picture of the Week

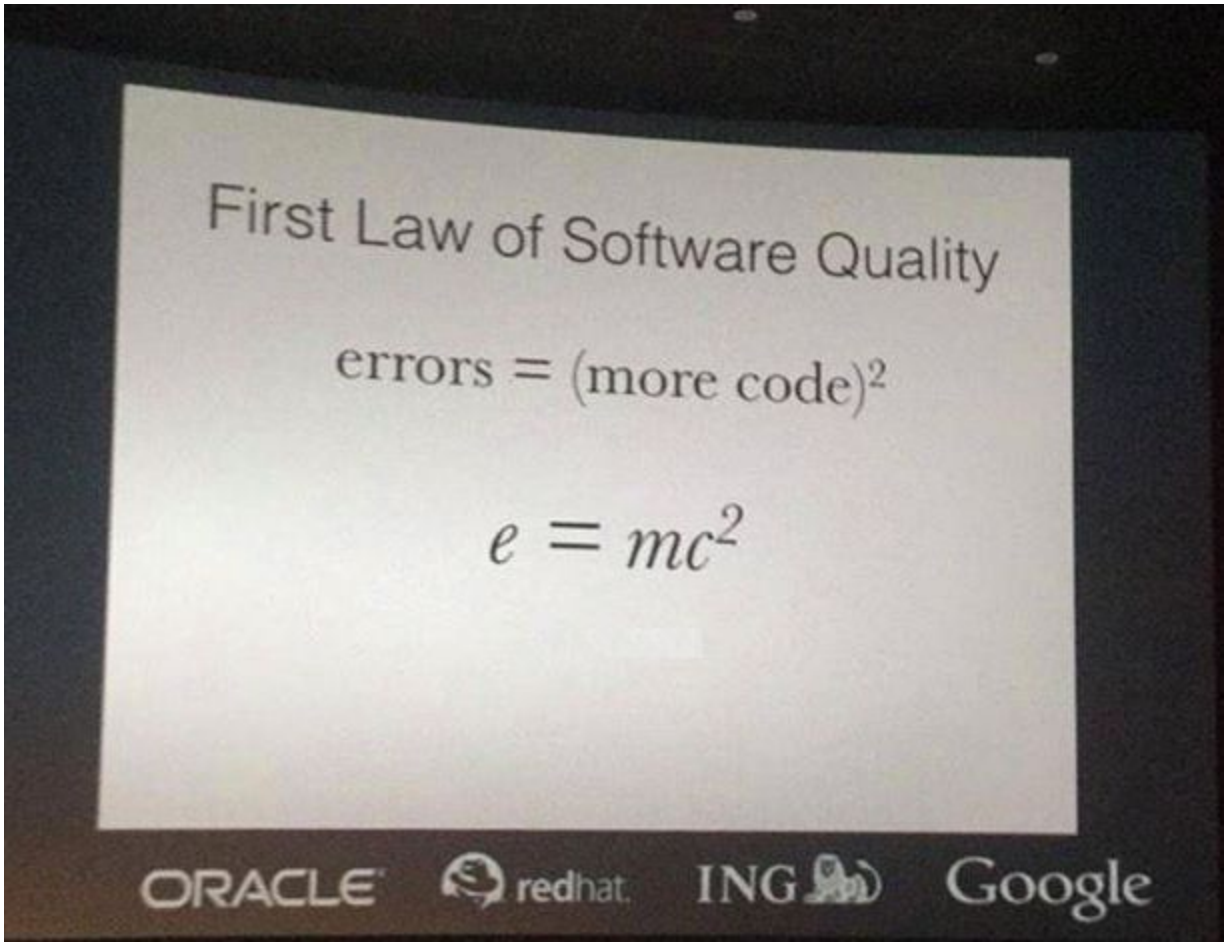

Security News

When NOT to upload a Work-in-Progress to VirusTotal?

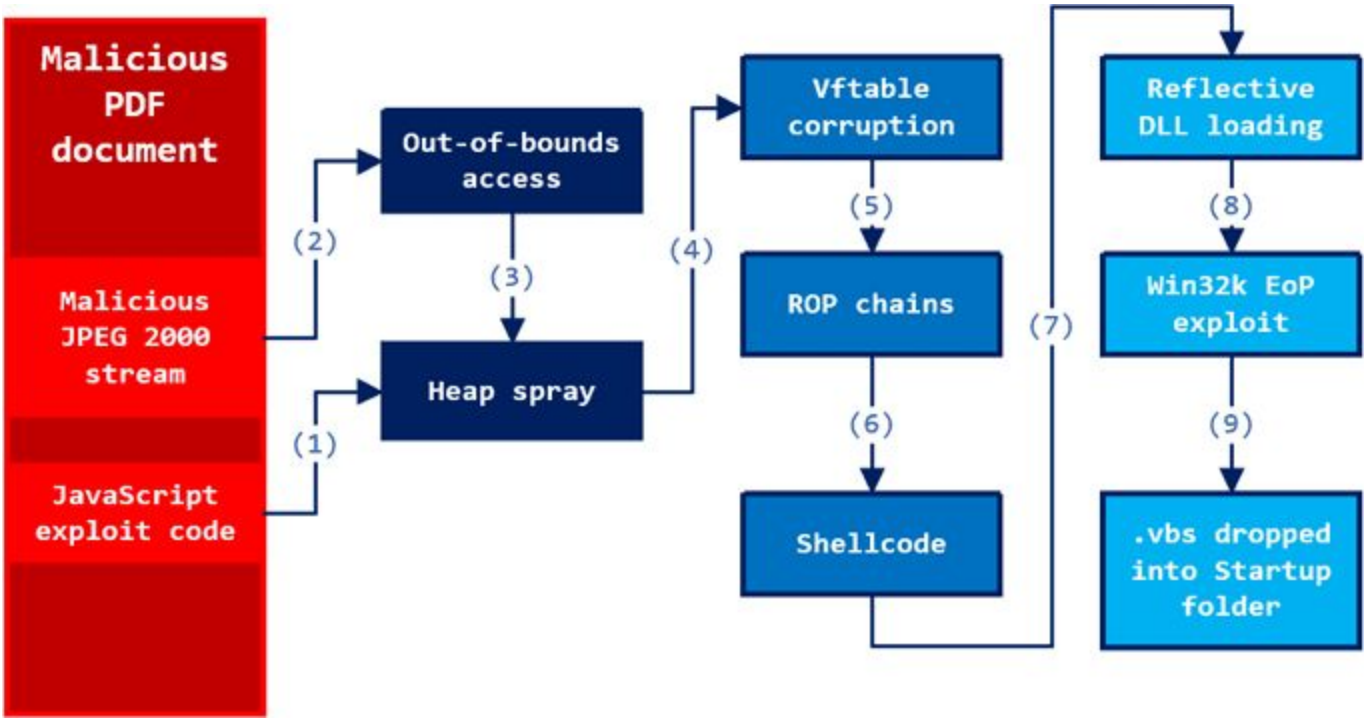

Two Zero-Day Exploits Found After Someone Uploaded 'Unarmed' PoC to VirusTotal
https://thehackernews.com/2018/07/windows-adobe-zero-exploit.html

What is “VirusTotal”?:

Headline: "Analyze suspicious files and URLs to detect types of malware, automatically share them with the security community."

"By submitting your file to VirusTotal you are asking VirusTotal to share your submission with the security community and agree to our Terms of Service and Privacy Policy. Learn more."

A 100% free and publicly available service where everyone shares for the benefit of everyone.

VirusTotal inspects items with over 70 antivirus scanners and URL/domain blacklisting services, in addition to a myriad of tools to extract signals from the studied content. Any user can select a file from their computer using their browser and send it to VirusTotal. VirusTotal offers a number of file submission methods, including the primary public web interface, desktop uploaders, browser extensions and a programmatic API. The web interface has the highest scanning priority among the publicly available submission methods. Submissions may be scripted in any programming language using the HTTP-based public API.

As with files, URLs can be submitted via several different means including the VirusTotal webpage, browser extensions and the API.

Upon submitting a file or URL basic results are shared with the submitter, and also between the examining partners, who use results to improve their own systems. As a result, by submitting files, URLs, domains, etc. to VirusTotal you are contributing to raise the global IT security level.

This core analysis is also the basis for several other features, including the VirusTotal Community: a network that allows users to comment on files and URLs and share notes with each other. VirusTotal can be useful in detecting malicious content and also in identifying false positives -- normal and harmless items detected as malicious by one or more scanners.

So, in other words, it brings social networking and sharing to the anti-malware industry, and everyone wins. As a single point of submission to, and evaluation by, ALL of the industry's A/V engines, new trouble can be found more quickly by end-users uploading something they suspect. But what's submitted also becomes available to the A/V community.

Back in March, ESET, one of the many A-V partners, was closely examining one of the many submissions to VirusTotal and raise an alarm when they thought they might have found a potential exploit for an unknown Windows kernel vulnerability hidden within the workings of a PDF submitted by someone to VirusTotal.

So... ESET forwarded their possible discovery to Microsoft, who confirmed the presence of TWO different previously known 0-day exploits in the PDF -- one on Adobe Acrobat & Reader, and the other a potent Elevation of Privilege (EoP) in the Windows kernel.

This is the story behind TWO of the patches that were released at the beginning of last month in May's patch Tuesday. Since sufficient time has passed for patched to be applied, the story can now be told.

According to the researchers, the malicious PDF was likely in the early development stage since the exploits were fully functional yet the PDF did not deliver a malicious payload and appeared to be proof-of-concept (PoC) code.

If this is the case, it would appear that the exploit's developers may have been curious whether their newly minted malware delivery system would be caught by any of the existing malware detection systems at the time.

Some of the reporting of this suggested that this may have been a mistake on the part of the malicious developers, suggesting that (quoting from some of the coverage) "It seems someone who could have combined both the zero-days to build an extremely powerful cyber weapon had unintentionally and mistakenly lost the game by uploading his/her under-development exploit to VirusTotal."

While this could be the case, Microsoft wrote that "The sample does not contain a final payload, which may suggest that it was caught during its development stages. Even though the sample does not contain a real malicious final payload, the author(s) demonstrated a high level of skills in vulnerability discovery and exploit writing."

Given a 60-day response window from VirusTotal upload to patch, this may also have been a carefully calculated gamble on the part of the attackers. If they were that good, it seems

unlikely that they would have made such a simple mistake. So they may have already been 100% ready to send a fully weaponized version of the exploit PDF to one or more intended targets, but, due to the nature of their target or targets, it might have been very important for them to judge the likelihood of detection, even if it meant possibly tipping off the A/V industry to two previously unknown and valuable vulnerabilities.

If the target was valuable enough and the probability of a successful spear phishing infection high enough, in an environment where raising an alarm by detection would be sufficiently detrimental, it's entirely conceivable that their upload to VirusTotal was coldly calculated.

Which strongly suggests that, as far as we know, someone, somewhere, was a victim, since VirusTotal had confirmed that for 60 days **not a single alarm** would be raised by any of the industry's anti-malware watchdogs.

Such is the world we live in today.

New problems for our LTE networks: The "aLTEr" attack

https://alter-attack.net/

Four security researchers have just released a research paper which pounds another nail in the well-nailed coffin of LTE. The paper will be delivered at the IEEE Symposium on Security & Privacy in 2019, but it's available today: "Breaking LTE on Layer Two"

https://alter-attack.net/media/breaking_lte_on_layer_two.pdf

As we know, the "Long Term Evolution" or LTE or 4G, is the successor to the "Global System for Mobile" which was GSM.

Abstract
Long Term Evolution (LTE) is the latest mobile communication standard and has a pivotal role in our information society: LTE combines performance goals with modern security mechanisms and serves casual use cases as well as critical infrastructure and public safety communications. Both scenarios are demanding towards a resilient and secure specification and implementation of LTE, as outages and open attack vectors potentially lead to severe risks. Previous work on LTE protocol security identified crucial attack vectors for both the physical (layer one) and network (layer three) layers. Data link layer (layer two) protocols, however, remain a blind spot in existing LTE security research.

About layers:
Physical (electrical) layer / RS-232, RJ45, 100Base-T, 802.11
Data link layer / Ethernet - packet collision handling, etc. / MAC addressing
Network layer / IP, ARP, ICMP
Transport layer / TCP, UDP, SSL, TLS
Session layer / TCP session establishment
Presentation layer / HTML, DOC, JPEG
Application layer / e-mail, SMTP, Telnet, FTP

In this paper, we present a comprehensive layer two security analysis and identify three attack vectors. These attacks impair the confidentiality and/or privacy of LTE communication. More specifically, we first present a passive identity mapping attack that matches volatile radio identities to longer lasting network identities, enabling us to identify users within a cell and serving as a stepping stone for follow-up attacks. We demonstrate how a passive attacker can abuse the resource allocation as a side channel to perform website fingerprinting that enables the attacker to learn the websites a user accessed.

This side-channel website fingerprinting is interesting research for its own sake, but it's sort of weak and obvious. In describing it they write: Meta-information on the data link layer leak information about the consumption of data per time unit. For example, if Bob watches a video, he uses more traffic compared to when he accesses a simple website. As a preparation step of the attack, Eve records popular websites and their layer two patterns. During the attack, she eavesdrops the meta-information and looks for similar patterns. In case she finds a match, she knows which website the victim visited -- with a certain probability.

So, yeah, if we can monitor traffic that we cannot penetrate, the pattern of traffic =could= potentially reveal something about the traffic's source. But this is a weak signal at best.

But their next revelation is potentially much more significant:

We present the ALTER attack that exploits the fact that LTE user data is encrypted in counter mode (AES-CTR) but not integrity protected, which allows us to modify the message payload. As a proof-of-concept demonstration, we show how an active attacker can redirect DNS requests and then perform a DNS spoofing attack. As a result, the user is redirected to a malicious website. Our experimental analysis demonstrates the real-world applicability of all three attacks and emphasizes the threat of open attack vectors on LTE layer two protocols.

The need for Encryption =AND= Authentication.

AES-CTR "Counter Mode" What is it. A bitstream. XOR.

When sending: Encrypt first then authenticate... So that upon receiving the payload is FIRST authenticated to verify it has not been modified and only then, if successful, is decryption performed. This prevents attackers from intercepting and modifying data to probe the decryption process.

But... LTE is an entirely non-authenticated protocol. AND it uses a simple XOR cipher which is trivially exploitable.

They write:

LTE uses mutual authentication on the layers above the data link layer to prevent Bob's phone from connecting to a fake network. However, the layers below are unprotected and an attacker can forward high-layer messages. The victim's phone still assumes that he is connected to the original network. For the user data redirection attack, we exploit that the user data is not integrity protected. Thus, an attacker can modify the content of a packet if they know the original plain text, even the packet is encrypted. In the case of DNS packets, we know the

destination address of the original DNS server. For the redirection, the attacker adds a specific offset, thus the DNS request is redirected to a DNS server under the adversary's control. The malicious DNS server, in turn, performs DNS spoofing to resolve the domain to a fake, malicious IP address. As a result, the phone sends a requests the wrong IP address.

Thus we have a fully feasible attack.

It is, however, an attack on the LTE link, which is either the Baseband processor (which we have ample reason to believe is not secure) or the radio link, where fake cell towers are known to exist. (Recall the sudden proliferation of "cell towers" in Las Vegas during DefCon and BlackHat.)

The team of three researchers from the Ruhr-University in Bochum, Germany and a researcher from New York University notified the relevant institutions such as the GSM Association (GSMA), 3rd Generation Partnership Project (3GPP), and telephone companies about the issues they discovered. But this is not a bug that can be patched. This is the current protocol which is present-by-design in every cell tower and 4G cellular device in existence. So it will not be fixed.

So what about the future?

The researchers are concerned that the LTE/4G "aLTEr" problem issue would also affect the upcoming version of the 5G standard in its current form because, although the 5G standard includes the data layer authentication which additional security features (stronger encryption at the data layer) to prevent aLTEr attacks, but these are currently optional.

And, in response to the attacks, the 3GPP group, which develops standards for the telecommunications industry, said that an update to the 5G specification might be complicated because carriers like Verizon and AT&T have already started implementing the 5G protocol.

And, what's more, the new "Diameter" system, which is the successor to the known-insecure SS7 system for managing inter-network handshaking, and is being deployed along with the new 5G protocol, is also fraught with problems. Although its design provides for the messaging encryption that was completely absent from the design from SS7, analysis of its actual implementation has revealed a large collection of problems and resulting vulnerabilities, rendering it no more secure, in practice, than its predecessor.

Researchers say that the Diameter misconfigurations they've spotted are in many cases unique to each network, but they usually repeat themselves to fall into five classes of attacks: (1) subscriber information disclosure, (2) network information disclosure, (3) subscriber traffic interception, (4) fraud, and (5) denial of service.

So what do we do?

We rely upon the upper layer: The encryption and endpoint authentication provided by TLS.

It's the job of all the various lower layers just to move the data from point A to point B. They perform, at best, a best effort to do so with integrity. But they cannot be relied upon.

EFF's next encryption initiative: STARTTLS Everywhere

Unless something bigger comes up next week, "StartTLS Everywhere" will be our topic and we'll cover eMail connection security and the EFF's latest initiative in detail.

In Edge, Chrome & Safari, the Spectre & Meltdown mitigations have been partially defeated.

Mitigations against these timing attacks have been incorporated into Firefox, Chrome, Chromium, V8, Webkit (Safari), and Edge and IE.

Though they vary by browser, in general the mitigations are:

1> Index masking of array objects
2> Site-Isolation feature in Chromium-based browsers
3> Disabling SharedArrayBuffer
4> Reducing precision of performance.now() timers
5> Adding jitter to the response of performance.now()

https://alephsecurity.com/2018/06/26/spectre-browser-query-cache/

"Overcoming (some) Spectre browser mitigations"

Because of this vulnerability discovery, browser vendors implemented different mitigations for this vulnerability. Some of them are meant to disable known methods for querying CPU cache state of memory slots (Javascript variables). These mitigations include the resolution reduction of the Javascript timer performance.now() and adding jitter to its results.

In our research we were able to overcome the cache access timing specific mitigations. Altough these mitigations cause a serious slowdown in our POC, they are not effective in preventing this attack. These mitigations also have some negative performance implications and are hurting the functionality of some legitimate Javascript web application use-cases.

UPnP-enhanced DDoS attacks are moving into the mainstream
https://www.bleepingcomputer.com/news/security/those-harder-to-mitigate-upnp-powered-ddos-attacks-are-becoming-a-reality/

Port based DDoD mitigation has been the norm.

DNS port 53, NTP port 123.

There was no ample source of servers on random ports.

But exposes UPnP services allow fully flexible port reassignment.

Unpatched WordPress Flaw Gives Attackers Full Control Over Sites
https://thehackernews.com/2018/06/wordpress-hacking.html
https://www.bleepingcomputer.com/news/security/unpatched-flaw-disclosed-in-wordpress-cms-core/

Coincidentally, we were just talking about path traversal bugs and it turns out that all versions of Wordpress have a known path traversal bug vulnerability... But it's one which apparently doesn't have the Wordpress authors very worried since they have known of it since last November.

Wordpress was notified of the problem last November by security researchers who made the discovery. But the Wordpress devs have, so far, not bothered to fix it.

The flaw was discovered in the PHP functions that deletes thumbnails for images uploaded to a WordPress site.

The researchers discovered that users who already have access to a site's posting post editor —and can thereby upload or delete images (and their thumbnails)— can insert code in a WordPress site that deletes crucial files which are part of the WordPress CMS core, something that should not be possible without access to the server's FTP.

So it's a very limited form of privilege elevation attack. It's not clear how useful deleting other files would be. It's been suggested that site could be hijacked by deleting the main wp-config.php, which is a site's config file. Attackers who deleted this file might be able to re-initiate the installation process and install the site using their own database settings, effectively hijacking the site to deliver custom or malicious content.

But... Any such attacker would need to first have other content-deletion rights on the site. And, dues to the requirement of an author-level account on a WordPress site, this limited vulnerability could not be mass exploited.

If there are large sites with large posting-privileged userbases who are worried about this, until an official fix is offered from Wordpress, the researchers have created a hotfix which can be found at the bottom of their report:

This hotfix is a bit of PHP code site owners can add to the functions.php file inside the site's currently active theme folder to prevent the

SpinRite

Ben in Greensboro, North Carolina
Subject: SpinRite recovers bad sectors but afterwards SMART status shows "good"?
Date: 27 Jun 2018 07:53:12
:
Hi Steve!

Long time SN listener and SpinRite user/abuser.

I recently came across a situation I found odd. A client's machine was having issues staying booted and BSODing. So, naturally, the first thing I did was run SpinRite at level 2 on the Windows partition and it it recovered SEVERAL bad sectors - which typically tells me that the drive is going bad. Of course, this fixed the trouble and the system then worked flawlessly.

But since the machine was still under warranty I figured that I might as wel get the drive replaced. So I contacted the manufacturer's tech support and while on the phone it passed their built in drive diagnostic, Windows WMI SMART check, and all tests showed the drive as "good"... Usually when SpinRite recovers sectors for a machine that won't boot it will fail any subsequent SMART checks. But not this time.

Do I have a fundamental misunderstanding of what recovery of a sector means?

I'm currently running a SECOND Level 2 scan on the whole C drive to verify it's operational before giving it back to the client and the ticket with the manufacturer temporarily still open (just in case I come across something).

An overview of the forthcoming (late next year) WPA3 upgrade

https://www.wi-fi.org/news-events/newsroom/wi-fi-alliance-introduces-wi-fi-certified-wpa3-security

The ill-fated previous work of the Wi-Fi Alliance is finally, after 15 years, receiving a long-needed upgrade. WPA2 arrived back in 2004.

Since devices will need to receive WPA3 certification, it's unclear what the firmware upgrade paths for SoHo routers will be. Presumably all of our Smartphones will be updated over-the-air.

However, this won't be happening tomorrow. WPA3 is not expected to appear until late next year, 2019.

Backward compatible, of course.

WPA3-Personal

WPA3-Personal brings better protections to individual users by providing more robust password-based authentication, even when users choose passwords that fall short of typical complexity recommendations. This capability is enabled through Simultaneous Authentication of Equals (SAE), which replaces Pre-shared Key (PSK) in WPA2-Personal. The technology is resistant to offline dictionary attacks where an adversary attempts to determine a network password by trying possible passwords without further network interaction.

(Remember that in WPA2, a passive attack on login credentials was found known as the Key Reinstallation Attack or "KRACK" attack. In this attack a passive listener can observe a user's login traffic, capture it and perform an offline brute-force attack on their password. Thus... as with any brute force attack, password strength is an important factor.)

Natural password selection: Allows users to choose passwords that are easier to remember

Ease of use: Delivers enhanced protections with no change to the way users connect to a network

Forward secrecy: Protects data traffic even if a password is compromised after the data was transmitted

WPA3-Enterprise

Enterprise, governments, and financial institutions have greater security with WPA3-Enterprise. WPA3-Enterprise builds upon WPA2 and ensures the consistent application of security protocols across the network.

WPA3-Enterprise also offers an optional mode using 192-bit minimum-strength security protocols and cryptographic tools to better protect sensitive data:

Authenticated encryption: 256-bit Galois/Counter Mode Protocol (GCMP-256)

Key derivation and confirmation: 384-bit Hashed Message Authentication Mode (HMAC) with Secure Hash Algorithm (HMAC-SHA384)

Key establishment and authentication: Elliptic Curve Diffie-Hellman (ECDH) exchange and Elliptic Curve Digital Signature Algorithm (ECDSA) using a 384-bit elliptic curve.

Robust management frame protection: 256-bit Broadcast/Multicast Integrity Protocol Galois Message Authentication Code (BIP-GMAC-256)

The 192-bit security mode offered by WPA3-Enterprise ensures the right combination of cryptographic tools are used and sets a consistent baseline of security within a WPA3 network.

Open Wi-Fi Networks

Wi-Fi Enhanced Open preserves the convenience open networks offer while reducing some of the risks associated with accessing an unsecured network by providing transparent unauthenticated data encryption to users. Based on Opportunistic Wireless Encryption (OWE) defined in the Internet Engineering Task Force (IETF) RFC8110 specification and the Wi-Fi Alliance Opportunistic Wireless Encryption Specification, Wi-Fi Enhanced Open.

The lack of authentication here is important since MITM attacks are possible. But at least passive sniffing of otherwise unencrypted traffic will finally be eliminated.

Wi-Fi Easy Connect reduces the complexity of connecting Wi-Fi devices with limited or no display interface such as IoT gadgets. Easy Connect enables users to securely add any device to a Wi-Fi network using another device with more UI, such as a smartphone, by simply scanning the product's QR code.

They state that public key crypto is used, so presumably the QR codes are per-device public keys matching an internal, secret and well-protected private key. This inherently creates an "out of band" static one-way communication path. Both next-generation WiFi routers and end-point devices "publish" their public keys. This would allow any device scanning the public key to choose an ephemeral key and encrypt it under the published public key so that it can only be decrypted by the device's private key.

Any facility offering free public Wi-Fi could either leave their network "open" using the new opportunistic keying system, or they COULD post the QR code from their W-Fi router and allow users to scan it for an MITM-proof wireless connection.