by Richard Moulds

November 17, 2017

ROCA, the role of key generation and decrypting of private keys

https://www.scmagazineuk.com/roca-the-role-of-key-generation-and-decrypting-of-private-keys/article/704343/

Richard Moulds takes a look behind recent crypto vulnerability headlines - the ability to calculate the private key of an RSA keypair purely by knowing the public key - and asks if they are a prelude to a 'cryptoapocalypse'.

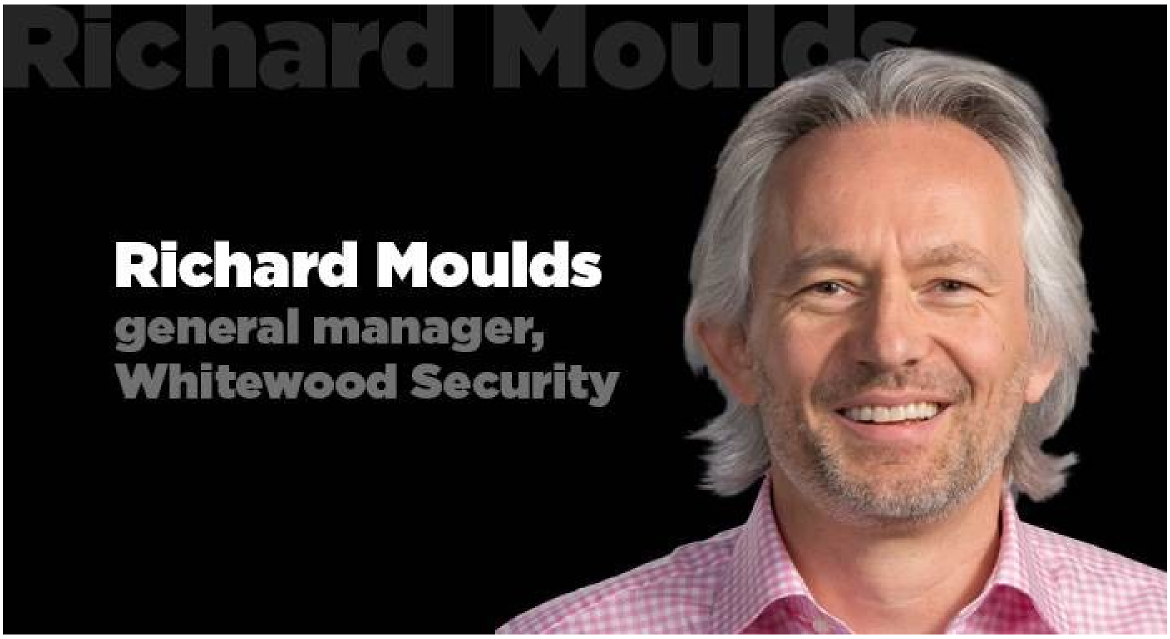

ROCA, the role of key generation and decrypting of private keys

It's been a busy time for crypto vulnerability stories. First there was the Key Reinstallation AttaCK (KRACK) that showed how a WiFi man-in-the-middle could trick WPA2 protocol handshakes into reusing encryption keys that are already known to the attacker. KRACK is significant because it points to a longstanding flaw in the WPA2 standard rather than an isolated implementation error - which means it could affect virtually every WiFi connected device. But perhaps even more worrying is the ROCA vulnerability, which has even wider ramifications and might yet serve as a dress rehearsal for the arrival of quantum computers.

The ‘Return of Coppersmith's Attack' (ROCA) makes it possible for attackers to simply calculate the private key of an RSA keypair purely by knowing the public key

– which is of course public, in the form of a certificate. The fundamental tenant of RSA asymmetric crypto is that determining the prime number factors of the public key is an immensely expensive task, way beyond practical computing reach. The problem is that it turns out that corners have been cut in the process of generating some of those keypairs, such that the factors can be found in just a few minutes for 1024 bit keys and a few weeks for 2048 bit keys. RSA 3072 and 4096 bit keys are still thought to be safe.

One of the big bottlenecks in generating RSA keypairs is that it takes a lot of CPU effort and therefore time, to find large random numbers that are prime. To overcome this, particularly in low power devices, algorithms have evolved to accelerate the process of checking whether a number is a prime number or not. It's one of these accelerator algorithms called Fast Prime that has been found to have a vulnerability that results in keys that can be easily factored. Fast Prime was used by Infineon and installed as a firmware library in various of their crypto hardware devices such as smart cards and Trusted Platform Module (TPM) chips.

It's an unfortunate irony that the organisations impacted most by ROCA are the ones that consciously tried to do key generation in the most secure way, in dedicated hardware. Anyone generating keys in software, for example, with OpenSSL, is fine.

The good news is that the vulnerability is easy to detect. But the bad news is that the vulnerability is also easy to detect. Normally, when key generation goes wrong, for example if there is insufficient entropy or randomness, to generate keys that are actually random, it can be hard for an opportunistic attacker to exploit the weakness, since vulnerable keys are indistinguishable from secure keys. The attacker has to spend a lot of effort just to find vulnerable systems before they even start to exploit them. In the case of ROCA it takes only milliseconds to determine if the certificate is weak. That's a big deal because weakness is something the attacker can and will, scan for.

The ROCA vulnerability is different from other scare stories and offers several important lessons –

1. **You don't always get what you pay for** - The victims in this story weren't trying to cut corners or caught tripping over their own bugs. Instead, they were trying to go the extra mile by investing in crypto-hardware. It's a double whammy; these organisations put this protection in place presumably because they have something

worth protecting, which is now at risk, compounded by the fact that the vulnerability can be easily spotted and targeted.

2. **Keys are a single point of failure** – Low-level tasks like entropy gathering and key generation are often taken for granted but are actually single points of failure that can bring down the whole crypto-house of cards. They go unmonitored and unmanaged but when they hit, they hit big. It's time to pay attention to your keys; knowing where they come from is as important as controlling how they are stored and managed.
3. **Embedded components have massive but unknown footprints** - Infineon chips are sold to equipment vendors. End-user organisations don't buy them or know if they have them. This means you might have to test every certificate for vulnerability. But you have way more certificates than you think and they're in places you might not expect or have access to.
4. **Remediation is easier said than done** – Although it's easy to test for the vulnerability, assuming you know how to find embedded certificates in systems you've never looked at before, updating firmware in proprietary devices and embedded systems will likely be a frustrating and expensive task. The situation is exacerbated by the fact that the devices and systems in question are designed to present a higher security posture and so have extra controls in place to prevent just the sort of changes that you need to make.
5. **The impact goes way beyond data theft; the infrastructure is at risk -** The ability to find private keys opens the potential to fake signatures and credentials, not just decrypt data. If you can fake code-signing signatures you can corrupt the infrastructure itself. Bad news for the IoT – Stuxnet for the masses.
6. **What value certifications?** - It's interesting to note that the Infineon chips that are affected were proudly marketed as FIPS 140 and Common Criteria certified. This raises obvious questions about the value of those particular certifications. We all know that certification schemes can often lag the actual market threat but in this case the vulnerabilities hit mature and ideally tightly scrutinised functions and yet they went undetected. When vendors routinely push labs to speed up the certification process and standards bodies entertain the idea of self-certification, are we missing the point?
7. **There's a right way to announce vulnerabilities** – It's tempting for anyone who discovers a vulnerability to immediately spill the beans and claim their 15 minutes of

glory. But there's a more responsible approach and in this case the researchers seem to have got it right.

a. Inform the vendors, and only the vendors
b. Give them a deadline with enough time to create a patch
c. Only go public about the vulnerability once a patch exists or if the vendors have ignored the issue, as a last resort
d. Give end users the tools to assess the risk and a realistic window to deploy the update; in this case,https://keychest.net/roca
e. Hold off from explaining the details of the exploit until most end users have had a chance to fix the issue

Looking further to the future, the ROCA vulnerability is eerily familiar to anyone who's tracking the threat posed by quantum computers. Sometimes called the ‘cryptapocalypse', it is expected that quantum computers will be able to execute Shor's algorithm, something that regular computers thankfully can't do. Shor's algorithm enables a private key to be calculated from a public key – sound familiar? Painful as it will be, the ROCA vulnerability is nowhere near as far reaching as the quantum threat but it serves to illustrate the issue. By way of contrast, the quantum threat is thought to impact all common asymmetric algorithms, not just RSA. It also applies to every device or application and isn't limited to specific chips or implementations. Worse still, quantum resistance might not be achieved through a simple software upgrade. The good news is that the quantum threat hasn't materialised, yet.

Back to the here and now and what the ROCA vulnerability shows us is that crypto isn't just about the algorithms. It really is all about the keys and in this case, how they are generated. Key generation represents a single point of failure and the failure is likely to be absolute, once a key is broken the game is up, trust is lost. As crypto becomes ubiquitous in securing the internet, clouds, mobile and the IoT we can't just take key generation for granted. We've historically judged crypto strength by how long the keys are. Maybe it's time to start asking how good the keys are and not just how many bits they contain. Anything less than true randomness is a risk, and the cost of failure can be immense.

Contributed by Richard Moulds, general manager, Whitewood Security