Eavesdrop resistant into the future

The TLS 1.3 protocol will provide the Internet with a new standard for encrypted communications. Led by Professor Marc Fischlin, a team of researchers at the TU Darmstadt collaborated in the analysis of the protocol and tested its cryptographic processes. These are more efficient and less error prone, but are they future proof?

By Boris Hänßler

Heartbleed, Triple Handshake, Crime – these cryptic-sounding names all refer to attacks on one of the core elements of the modern Internet: the Transport Layer Security (TLS) protocol for encrypted communications. We all use it to protect our data whenever we google something, buy from online shops or send an email. Generally speaking, we can rely on the TLS protocol: the worst hacker attacks carried out in the last few years were all the result of failures to properly implement it. Nevertheless, many experts were still unhappy with the protocol. On the one hand, researchers themselves have periodically been discovering minor security breaches in the protocol itself. On the other hand, many Internet-based companies have been finding the protocol too cumbersome and not fast enough.

"The reduction of round trips is among the most important innovations, but also entails certain risks."

That's why the international "Internet Engineering Task Force" (IETF) has introduced a new protocol known as TLS 1.3. Together with his team, Marc Fischlin, Professor of Computer Sciences at the TU Darmstadt, used their expertise in cryptography to carry out an analysis of the new standard. During the development process, which spanned several years, they tested the proposed procedures and their new functions to ensure that they really are sufficiently secure and that they will be able to keep pace with future technological developments.

One security risk that Fischlin's team has been looking into is the result of fewer so-called "round trips" in the new protocol. To establish a TLS connection, the client – for example the PC used by a customer of an online shop – and the shop server negotiate an encryption key in several steps. The contact is initiated by the customer's PC, which essentially says: "Hello. I'd like to communicate with you and propose the following encryption keys." The server then selects one of the available encryption keys and simultaneously transmits an official authentication certificate, which confirms that it really is the server of the shop in questions. The customer's PC accepts the key in turn and, following a number of additional steps, both parties declare the negotiation processes to be complete. Only then can the actual data exchange take place. This technical dialogue consists of six steps in total. One objective of the new protocol was to reduce this negotiation process to just four steps, which the developers succeeded in doing by combining two formerly separate processes.

If the client and server already know one another, for example because the customer has already bought things from the online shop, then the new TLS protocol permits them to communicate immediately. Computer scientists refer to this as a "zero-round-trip-process" because no additional encryption key negotiation rounds are required. The customer's PC authenticates itself using a so-called "session ticket", which it receives and saves during the initial encryption negations upon first contact. It can use the ticket to transfer its relevant data.

The reduced number of round trips is not something that we would notice in our everyday Internet usage. Even the old version of the protocol was so fast that its impact on our communication speeds was negligible. However, the situation is entirely different for search engines such as Google, which, at the last count, registered some two billion search queries per year. Every time someone reloads the search engine in their browsers, a new encryption key has to be negotiated. Therefore, the round trip reduction in the updated protocol represents a significant traffic load reduction for companies such as Google. "Without doubt" Fischlin explains, "this is one of the most important innovations of TLS 1.3. But it does involve certain risks".

One of these risks is the subject of a current research project, which Fischlin and his team will be presenting at the 2nd IEEE European Symposium on Security and Privacy in Paris in April. The subject of the study are

Information

Institute of Computer Sciences and Complex Cryptography
Prof. Dr. Marc Fischlin
marc.fischlin@cryptoplexity.de
http://bit.ly/2mlTbs9

so-called "replay attacks" in zero round-trip scenarios. In the course of such an attack, the hacker would attempt to intercept the session ticket. Whilst he or she would not be able to use it to read or alter the transmitted data, they could use it to send multiple requests to a given server. If, for example, a user were to order a book from an online retailer, the hacker would be able to reorder the book thousands of times at the user's expense.

For this reason, the TLS 1.3 protocol should enable online shops and other service providers to check if user requests are repeats of earlier requests; whether, for example, the same product is being ordered multiple times. If so, then the session ticket will be invalidated, and the server and client would have to negotiate a new encryption key, thus taking the would be hacker out of the loop. Fischlin and his team were able to prove that the new protocol meets the new security requirements even in the face of any conceivable exceptional case. "There are still residual risks involved", says Fischlin, "but we consider them to be so minuscule that we consider the new protocol to be robust".

The Darmstadt researchers are also looking ahead to future functions for TLS 1.3. Anti virus software manufacturers, for example, would like it if their software could search encrypted data, instead of having to wait until the files are decrypted. "We're working on procedures that would enable this", says Fischlin: "One option, for instance, would be to carry out computations directly on the encrypted data string. The virus scanners would be able to recognise damaged or compromised code based on the encryption pattern. Of course, we'd have to ensure that standardised information included in the files, such as banking data, could be read out using the same method."

In the distant future, cryptography experts will be faced with another problem if quantum computers are ever realised, which would render current encryption processes, based on the so-called Diffie–Hellman key exchange method, obsolete. This process is deemed secure because it is based on a mathematical problem that is fundamentally intractable for current computers. Quantum computers, by contrast, could solve it. As Fischer explains: "No one can say whether such computers will ever be realised. But, as soon as they are available, all these procedures would become obsolete overnight. We need to be prepared for this."

One potential replacement would be the so-called learning with errors problem (LWE), which, it is assumed, could not be solved by quantum computers. So why isn't it being used already? "Ah", Fischlin explains with a grin, "that would increase protocol latency again, which just goes to show that cryptographers are not going to run out of research subjects any time soon!"

The author is a science writer.

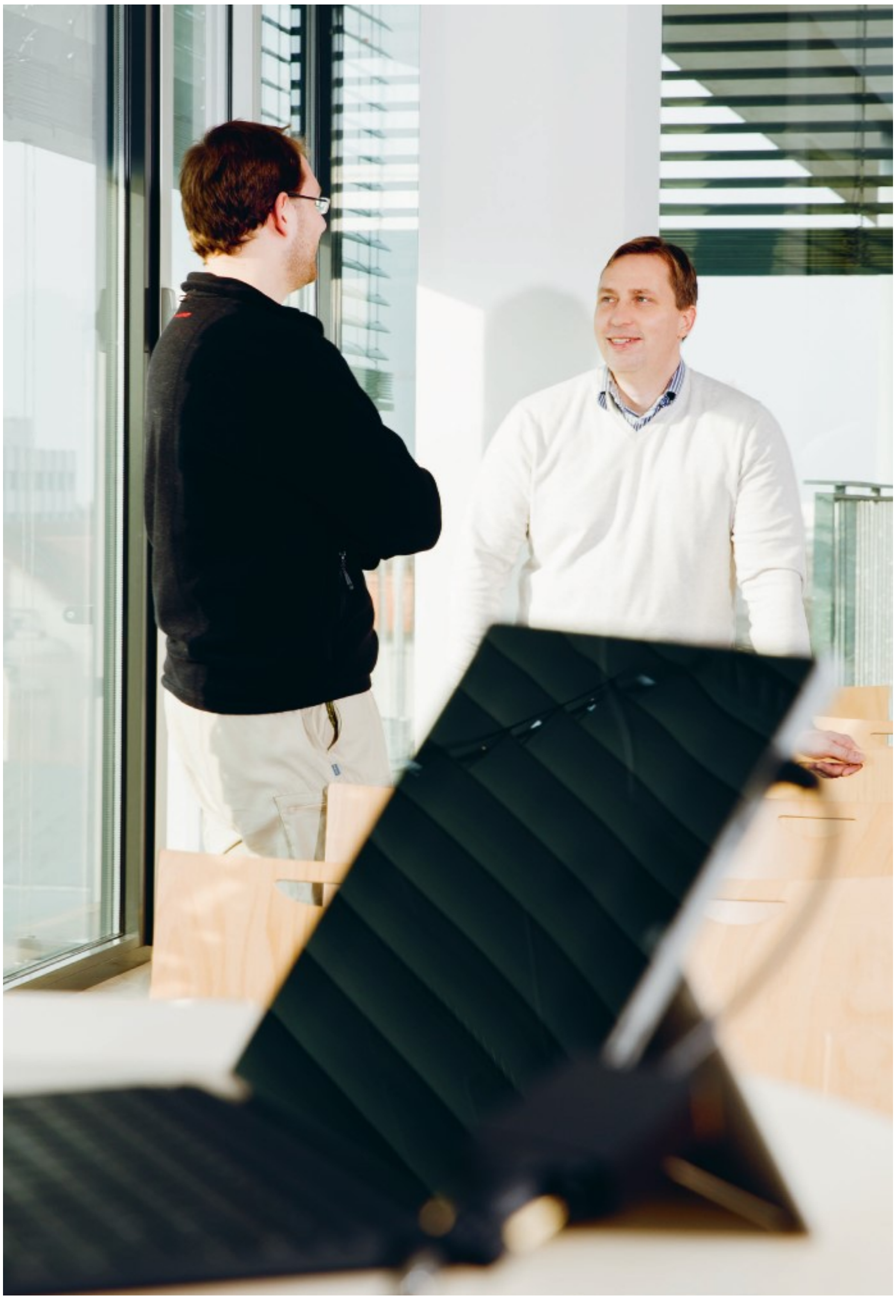

Photo: Katrin Binner

Carrying out research into cryptography, security and complexity theory: Professor Marc Fischlin (ed.)