Replay Attacks

Author: csirt@bit.admin.ch & outreach@govcert.ch
Topic: Replay Attacks
Date: 14th of June 2020
Classification: **TLP WHITE**

Content

Introduction

During the public security test, various testers mentioned that the risk of replay attacks exists and could pose a serious threat to the application as they may be used to poison the system with false information and thus either put non-affected persons into quarantine or spread fear with higher than real infection rates.

We have mentioned the risk of replay attacks as well as the possibility of eavesdropping in our report[1] as well:

Replay attack is the only real sabotage possibility we could find in the protocol: An attacker can collect EphIDs of people with a high probability of future positives with a very sensitive receiver, e.g., near a drive-in test center or a hospital in general, send these via internet to a very different location where a lot of non-infected people are expected (like in residential areas), and replay them there using a very strong Bluetooth signal. This would cause a lot of false detections.

We would like to shed a bit light on these types of attacks in order to show the actual threat that may originate from this type of attack.

Due to how SwissCovid works these problems cannot completely avoided as a BLE beacon has no authenticity by itself as this might lead to other attacks against the privacy and anonymity of its users. There are however a few approaches that would reduce the risk and increase the price an attacker has to pay for a successful attack.

Various reporters presented detailed scenarios for such attacks; we try to summarize them as follows:

Replay Attacks

One attack category deals with replay attacks in order to generate false alerts about infections:

- An attacker can collect a large amount of EphIDs by sniffing the BLE beaconing at a location with a higher probability of infections, e.g. near a hospital or a testcenter
- He can then take the sniffed EphIDs to another place with a high visitor frequency (e.g. a train station) and send these EphIDs with a higher sending power in order to make the other smartphones believe that they were near to the sender for an extended period of time (currently 15 minutes)
- This is going to lead to false alerts about possible infections for users that did record these EphIDs.

[1] https://www.melani.admin.ch/dam/melani/en/dokumente/2020/riskestimationproximitytracingappendix.pdf.download.pdf/Risk-Estimation-Proximity-Tracing_Appendix_Signed.pdf

Unmasking users by eavesdropping EphIDs

Another issue raised is the possibility to demask users who uploaded their keys after a positive test by preceding eavesdropping EphIDs. We acknowledged in our report[2] that this attack was feasible as well:

An attacker who is in proximity of the victim can eavesdrop the EphIDs. Currently, no defense against that is implemented. EPFL proposed spreading the EphID across low-energy beacons using k-out-of-n secret sharing. However, this is not implemented, and it is not clear if this can be done using the Google/Apple APIs in the future.

As one researcher pointed out, there exists an additional attack vector by using existing collection of BLE beacons where an attacker might combine these with information gained elsewhere. There are several organizations that collect Bluetooth beacons in order to geolocate users. One good example are Facebooks location tracking or several SDKs that are being used for advertisements.

Eavesdropping EphIDs is not possible on the most recent versions of Apple API, only on Android. Apple sets the EphID to 0 if any app tries to access BLE beacons used for proximity tracing, while Google allows accessing it. We believe that Apple's approach by hiding the actual EphID from Apps is the best way to go and we would welcome if Google did the same and implemented a filter. This would reduce the likelihood of easily available side-databases to some extent, but not completely. Companies scanning for BLE beacons such as WIFI Hotspot operators or payment providers might still be able to collect enough BLE beacons on different hardware and use additional information they gain to demask the user. Also, even Apple devices running old iOS versions can be used in this sense.
The public should be informed that people can turn on and off the app at any time and so stop broadcasting EphIDs for defined periods of time. It is important to keep the app running whenever infection situations with unknown people can occur, but it is better to turn it off at home, which reduces the replay attack risk on the receiving side, when in places that should later not be exposed, or when at work if a risk of BLE collectors operated by the employer exists. Using the app is not a binary decision, but can be adapted by users depending of their current environment.

It is important to notice though that the privacy risk only affects diagnosed people, i.e. those that received a positive test result and uploaded their TEKs subsequently, and not at-risk (i.e. warned) people, as claimed in by one researcher. The fact the number of infected people is much lower than the overall number of users or even of at-risk users shows that the attack surface is quite small and restricted to patients who will need to go into isolation by law anyway, which poses a much larger impact to their privacy than a theoretical risk due to preceding eavesdropping. Also, the time range where this risk exists for these users is restricted to the contagious window, usually a few days.

Another and maybe even more serious threat could originate from organizations that operate stationary tracking systems such as payment providers relying on Bluetooth or WIFI operators. There is no real safeguard with the current design against this attack vector, however there are only few operators of such systems and they are under the Swiss jurisdiction which gives at least some protection on the legal level.

[2] https://www.melani.admin.ch/dam/melani/en/dokumente/2020/riskestimationproximitytracingappendix.pdf.download.pdf/Risk-Estimation-Proximity-Tracing_Appendix_Signed.pdf

Effect on the privacy of the users

We believe that under normal circumstances, the privacy of the users does not constitute an unacceptable higher risk when using the app. If a user has a smartphone with Bluetooth enabled (e.g. for headphones), she accepts certain risks associated with this technology. The same is true for the Swisscoovid App. One might argue that the overall attack surface for the population rises because users are pushed into activating Bluetooth. While this is true, we believe that many people already have Bluetooth enabled and that Bluetooth based proximity tracing is still the better option than using actual geolocation information. We do not see any other better technologies that could be made ready within the given timeframes.

Effect on the data quality of the system

It is difficult to estimate the likelihood that someone really tries to poison the whole system by large-scale replay attacks or to try to put individuals wrongly into quarantine by a targeted replay attack against some persons. As notified individuals are now eligible by law to free testing, targeted attacks would only be effective for a short time (until the results of the tests are known). Large-scale replay attacks would generate a high number of tests and thus be easy to detect. In case the worst case, the app could then be disabled. The damage would be limited to development costs of the app, and a lost opportunity. The alternative would be not do offer a proximity tracing app at all – even an app using a centralized approach would be vulnerable to replay attacks to a certain, though lower, degree.

In the end, this is a question that cannot be answered on a technical level alone. While such an attack is possible, it is associated with high costs for the attackers. As such, we believe that the benefit of the app is higher than the potential risk of such an attack.

Possible countermeasures

It is difficult to effectively combat these scenarios. One possibility would be to use some kind of coarse location information that is combined with the EphIDs. By doing so, the geographical space where these EphIDs are considered valid can be controlled. Such a coarse location information could be XOR-ed with the (daily) rolling proximity identifier key before AES is applied to calculate the rolling proximity identifier (10-minute interval), and the smartphone would have to memorize them locally together with timestamps. When at a later point TEKs are uploaded, only people with identical location information codes at the same time could apply the same XOR and so find a match. These location information codes can also be regularly changed and need not contain actual location information; it's only relevant that people in a close geographical area can somehow share this data.
We suggested such an approach, but it was declined because it may raise privacy concerns as the app needs to access alternate geolocation information, such as GPS or GSM antenna information. However, Google and Apple could add such a feature into their API in a well-documented manner that would not require asking the user to grant access for this type of information, as long as this is part of the exposure notification API in a transparent manner. This decision would be up to Google and Apple though, and so not within scope of the app code. It still might raise concerns, and increases the complexity of the implementation.

Another possibility to at least have some chance of anomaly detection would be the use of stationary measurement points that would be able to detect the flooding of EphIDs originating from single senders. This however has also privacy drawbacks, at least to some extent, and the detection of such anomalies would need a lot of research, testing and continuous monitoring.

The best way to reduce the effect of eavesdropping EphIDs is the proposal by EPFL to spread the EphiDs using k-out-of-n secret sharing.

It may be possible to detect such anomalies later on when people are tested after having gotten an alert and if a contact tracing does fail to show any possible infection candidate. This however is beyond the possibility of a technical team to estimate. We believe that the most important thing to do is to accept that there are residual risks and to perceive the app as just one additional data source for the handling of the pandemic.

incidents@
govcert.ch

Digital unterschrieben
von incidents@govcert.ch
Datum: 2020.06.15
09:09:49 +02'00'