Bachelor thesis
Computer Science

Radboud University

Breaking the chain: Security analysis of a pilot system for detainee monitoring in the Dutch Custodial Institutions Agency

Author:
Ward Maxime Theunisse

First supervisor/assessor:
dr. Peter Schwabe
p.schwabe@cs.ru.nl

Daily supervisor:
dr. Veelasha Moonsamy
v.moonsamy@cs.ru.nl

August 27, 2017

This page intentionally left blank

Abstract

Since late 2016 a pilot project to develop and test an IoT-based detainee monitoring system has been going on at a small-scale facility of the Custodial Institutions Agency of the Netherlands (Dienst Justitiële Inrichtingen (DJI)). This pilot system allows detainees to go on leave without accompaniment, while their location is being tracked. In this thesis we investigated to what attack scenarios such a system is subject and/or vulnerable and how these attacks can be prevented or have their impact reduced. We also looked at possible attacker goals and the cost and feasibility of attacks. This thesis is highly relevant in a societal sense, because the use of this pilot system could conceivably be or be not expanded within the DJI based on its assessment, at which point the pilot system would need to be trustworthy. During our research we found certain attacks that break the security of the pilot system in its current state, along with countermeasures against these attacks.

Contents

1. Introduction

In a contract with the Dutch government and in cooperation with the Custodial Institutions Agency (DJI) a pilot system has been in development to monitor the location of detainees when they go on leave. This project currently sees active experimental use at a small-scale location of a youth detention centre. The project was originally intended to run at the penitentiary of Zoetermeer, but it was moved to the location of another experiment of DJI: a small scale low-security youth detention annex. This happened because the prison it was supposed to run at closed down. Its use may expand after this trial. It is very important that a security analysis of this pilot system is made before it replaces existing structures, because potential breaks of the pilot system could allow detainees to not follow the agreed-upon conditions of their leave or even attempt to run away while fooling the centre into believing they have not. This analysis intends to allow the involved parties to better asses the risks and improve the security of this pilot system. In this research:

- We propose an attacker model with respect to detainees in conventional (youth) detention centres against monitoring systems such as the pilot system we are working with. (Section 3.1)
- We propose certain practical problems that may interfere with the pilot system (Section 3.2)
- We create an overview of attacks on this pilot system, constructed as a result of sniffing of communications and inspection of decompiled sources and documentation (Section 3.3 and 3.4). For each of these attacks we discuss:
  - the method and prerequisites of the attack
  - the consequences of the attack being successful
  - what the risk of this attack occurring is
  - how to hinder or prevent this attack or reduce the impact of it being successful

By doing this we want to answer the following question: To what attack scenarios is a system such as the one tested at the DJI subject and/or vulnerable and how can these attacks be prevented or have their success impact reduced?

Disclaimer:

In order to be able to perform our research for this thesis, we have received materials from S & T in the form of hardware and passwords. As part of this agreement we consistently refer to the pilot system developed by S & T as a "pilot system" instead of a "system" even if it might hinder readability in places. This is to stress that this pilot system is not a "product or even concept-product". We have received no monetary compensation for this research. All findings of this research have been shared with S & T at least 3 months before publication.

2. Preliminaries

2.1 Context of the project

Since 2012 the Ministry of Security and Justice, the Ministry of Defense and the National Police annually organise the so-called "Security Innovation Competition"[1] in order to find innovative ideas around a given theme or challenge that could be of use to the Ministry. The winner of this competition wins a contract of two hundred thousand euro to develop the concept in cooperation with the Ministry. In 2015, when the theme was "Smart access control", the winner of this competition was the so-called "Cardio Access Key"[2] by Science & Technology Corporation[3] (S & T).

Their idea made use of a third-party product called the Nymi Band; A bracelet authenticated with heart rate that could be used to grant access (authenticate) as long as it was not removed from the wrist or cut (more about this this in section 2.3.1). Originally the idea was to experiment with this product at Dutch prison as a way to regulate access to government buildings.[4] During the process of development, experimentation and cooperation with the Custodial Institutions Agency, the goal of the project has pivoted into monitoring detainees on leave who might not need accompaniment, but where being able to track them would still be beneficial.

The project was intended to take place at the prison in Zoetermeer, but because this prison closed down, the project moved to the small scale youth detention facility which it was at at the time this thesis was written. This small scale facility is itself an experiment of the DJI to test a lower level of security for detainees that are considered low risk. Within the DJI, the small scale facility falls under the JJI (Judicial Youth Detention Centres). No tracking system has ever been utilised within the JJI until this pilot. During the continued development this newly proposed monitoring system is under experimental use in the tracking of juvenile detainees on leave as an addition to the internal pilot of the DJI.

The pilot will end in the winter of 2017.

[1] https://www.defensie.nl/onderwerpen/innovatie/veiligheid-innovatie-competitie

[2] https://www.defensie.nl/actueel/nieuws/2015/10/16/%E2%80%98hartslagsleutel%E2%80%99-wint-innovatiecompetitie

[3] https://www.stcorp.nl/

[4] http://nieuws.securitas.nl/2015/11/30/cardio-access-key/

2.2 General overview of the pilot system

A high-level description of the pilot system (figure 1) as it is in place currently can be described as follows: A detainee who goes on leave gets fitted with a Nymi band by a detention group leader (in the context of penitentiaries this would have to be another employee). The group leader brings this Nymi band into authenticated status with a specific phone which stays in the center (the Nymi Companion App phone, or NCA-phone). The band then communicates its status with a coupled locked Android phone that the detainee has to carry with them while on leave. This phone periodically communicates its location, battery status and when it has last seen the (authenticated) Nymi band back to the S & T servers using the Dutch cellphone network. The employees of the detention center are able to access a webinterface which receives these data from the company servers. The interface displays the location of all active "rooms". A "room" is a user identifier referring to the room number of the detainee on leave that is tracked by the pilot system. Employees link a Nymi band-phone pair to such a "room" using a desktop helper application by communicating the wanted pairings to the company servers. In addition to the current location of a "room" (i.e. an active locked Android phone running the tracking app, which is paired with a Nymi band), the webinterface also shows the time the S & T servers received the last update message from the locked phone (heartbeat) and plays an alarm if this has been too long ago. Furthermore, the webinterface also displays the battery status of devices and is able to show the location history and warning history for specific "rooms".

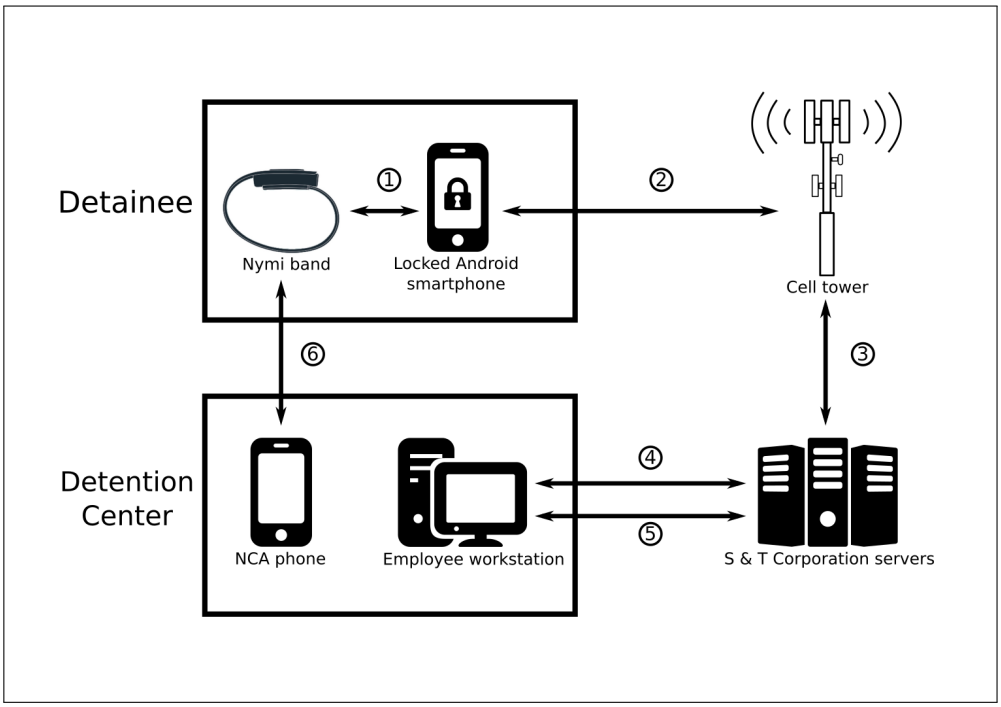

Figure 1: A schematic overview of the pilot system

In the next section, we provide a further description of the individual components of the pilot system. After this section, we will explain how leave works within the context of conventional youth detention and how what might become an escape attempt is handled.

2.3 Components of the pilot system

2.3.1 Nymi band

The Nymi band[5] is a device designed to be used in multi-factor authentication and is worn around the wrist. Hardware-wise, the band consists of the following relevant components:

- ARM Cortex M4 Processor with 2 Mb flash memory
- Bluetooth 4.1 module
- NFC module
- Accelerometer / Gyroscope
- Secure element (and cryptographic processor)
  According to the manufacturer, the secure element is an NXP chip capable of SHA-256, NIST P-256 and Secp256k1.
- Two single lead dry electrodes; one on the inside of the band, one on the outside. These are used for the purpose of Electrocardiogram-based authentication.
- A continuity circuit, closed by clasping the device, making contact with the connector pins on the top of the Nymi band, as depicted in figure 2.

The Nymi band works as follows:
The device can be in either an authenticated state or not. When the continuity circuit loses power by being unclasped, any authenticated status is lost. The band will only communicate with coupled applications when it is the authenticated status. In the context of this project, the coupled application is the tracking app on the locked Android phone. The Nymi Companion App can always communicate with the Nymi band. The band can be brought into the authenticated status using the phone which has the Nymi Companion Application that is paired with this specific Nymi band (or, of course, has the same keys) using either ECG readings or a password.

In the pilot project at the DJI, the beta version of the Nymi band and its Software Development Kit is used.

[5] `https://nymi.com/`

Figure 2: What a Nymi band looks like.
(Copyright Nymi Inc. 2017, `www.nymi.com`, usage protected under Fair Use)

2.3.2 NCA Phone

Along with a Nymi band, a phone (either iOS or Android) with the so-called "Nymi Companion Application" (NCA) is used. This application serves as an interface through which a Nymi band can get into an authenticated status. Upon first setup of the device, it performs a Diffie-Hellman key exchange with the NCA phone and provides it with ECG readings. On the NCA phone, a profile is created for this Nymi band using processed ECG readings and a chosen password, i.e. the generated session key is linked to the password and ECG data. The Nymi band does not store this information. If you want to change the saved ECG data to be able to authenticate another person, or just change the password, you need to have the band in the authenticated state. To get the band in authenticated status, you can either have the NCA phone request ECG readings from the band, or enter the password on the phone when it is within communication range. The firmware of an authenticated Nymi band can also be updated using the NCA. Back in 2015, the company producing the Nymi bands have expressed that a future update to the Nymi firmware would include the option to authenticate on the band itself *(p. 6, Nymi Inc. (2015) [1])*, making the using of an NCA optional. As of late 2016 a version of the product that supports this option exists, but for fairly obvious reasons this feature should not be available to detainees (in the form of biometric authentication).

In the application of the Nymi band in the pilot at a juvenile detention center the NCA Phone stays inside of the facility. When a detainee goes on leave, his device is brought in authenticated status by one of the supervisors of the detention group that detainee is a part of. If the system were to be used in other institutions within the DJI, outside of juvenile detention centres and forensic psychiatric centres, the system of "groups" (such as detention groups) is not used, so another type of employee would need to fulfil that role there. In prisons, this task would maybe fall to the "badmeester" (the employee responsible for visitation and search, entry and release). It is imperative that the detainee himself cannot authenticate the device, because

this would allow him to subvert the monitoring process by reauthenticating the band without his arm in it.

2.3.3 Band - Locked Phone Communication

The primary communication channel that the Nymi band uses to communicate with phones is Bluetooth Low Energy with an added layer of asymmetric crypto after the key exchange. This layer is an undocumented proprietary protocol, likely using either NIST P-256 or Secp256k1, given that those curves are supported by the crypto chip on the band. An authenticated band can communicate with phones runnning Nymi Enabled Applications (NEA) (i.e. apps using the Nymi Software Development Kit). In the case of the pilot system, this NEA is the tracking app. Upon first discovery of a Nymi band by an NEA, the NEA initiates a Diffie-Hellman key exchange (just like with the NCA) in which public-private keypairs are generated specific to that pairing of band and application. These keys are stored on the phone, using OS-specific security features.

In the case of the pilot project at DJI, the keys of the tracking app are stored under the `/data/user/0/` directory, which is the data directory for an Android root user. Interestingly enough, the Nymi API doesn't use the Android Keystore feature, which is a keyring system specifically designed to make it more difficult to extract key material from the device.

2.3.4 Locked Android phone

In the monitoring pilot system, the detainee on leave is provided with a locked Android phone running the tracking app. Locked in this context means that the detainee does not know the passcode of the phone, which is needed to get past the lock screen. The employee of the facility links the app with the Nymi band that the detainee is fitted with (by letting the app perform the key exchange). The phone employs a whitelist, so it can only be called by specific phone numbers related to the detention center or S & T. This phone runs the tracking app that checks if the band is still authenticated and communicates status and location to the S & T servers. The band is not used to check if a pulse or movement still exists in the band; Only if it is still authenticated. (Not that checking this would add much extra security, because both are fakeable and authenticating is the hard part for an attacker. Checking the pulse against the heart profile would not be desirable because heart rate changes too much based on mental state, bodily exertion and posture leading to inconsistent biometric authentication.)

As an aside: detainees also carry their own cellphone on them during leave, and these are generally used if the detention center needs to contact them. The locked phone we were provided with and which is used in the pilot project was a ZFD VFD 600 (Vodaphone Smart Prime), running Android

6.0.1 with the May 1st 2016 patch level.

2.3.5 S & T server

The server runs Microsoft IIS 10.0. For earlier versions of IIS (6.0, 7.0 and 7.5) there are published remote code execution vulnerabilities (`CVE-2017-7269`, `CVE-2010-3972` and `CVE-2010-2730`), but for version 10 no such published vulnerabilities exist at the time this thesis was written. In terms of practical research, we have not focused on the server because we were asked not to for reasons of uptime and stability, seeing as the pilot system has to stay online because it is actively being used. The server presents the webinterface for the detention center and processes the heartbeat messages of the detainee phones.

2.3.6 Locked Phone - Server Communication

Every 60 seconds, the tracking app on the locked Android phone restarts. This restart is a workaround to free up process memory, because of an unplugged memory leak that is part of the beta version of the Nymi Software Development Kit. Upon restart, if the Nymi provision exists and the location accuracy available is sufficiently high, it tries to send a message using a `HTTP POST` over `TLS` to the hardcoded server address. This happens over the 4G Vodaphone network. In the `TLS` handshake, the server authenticates itself against the phone, but not vice versa. If it fails to set up the `TLS` connection or a bad certificate is provided, no message is sent. The messages that the phone sends are of the following form:

```
<NymiID>; + <Last known location>; <Last accuracy>; <Unused field>;
    <Time at which last location was recorded>; <Time at which the
    Nymi was last seen>; <Battery level>; <Password>;
```

The Nymi ID is a unique identifier of the Nymi band - Phone Application pair which consists of 32 alphanumeric characters and is generated alongside key information during the initial DH key exchange between the band and the phone. For the location and accuracy, the built in Android Location API is used. This API uses GPS. In theory some location information could also be obtained through the cell network provider because they know to which cell tower the phone is connected. The pilot system does not do this though. The last field of the message is a secret password. This password is a universal password that is the same for all phones running the application, and is hardcoded in the tracking app source code.

One of the first things to check when investigating the security of a `TLS` connection is look at the cipher suites that both parties support. The client application (the tracking app) supports the following cipher suites:

```
Cipher Suites (20 suites)
```

```
    Cipher Suite: TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256 (0xc02b)
    Cipher Suite: TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384 (0xc02c)
    Cipher Suite: TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 (0xc02f)
    Cipher Suite: TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 (0xc030)
    Cipher Suite: TLS_DHE_RSA_WITH_AES_128_GCM_SHA256 (0x009e)
    Cipher Suite: TLS_DHE_RSA_WITH_AES_256_GCM_SHA384 (0x009f)
    Cipher Suite: TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA (0xc009)
    Cipher Suite: TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA (0xc00a)
    Cipher Suite: TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA (0xc013)
    Cipher Suite: TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA (0xc014)
    Cipher Suite: TLS_DHE_RSA_WITH_AES_128_CBC_SHA (0x0033)
    Cipher Suite: TLS_DHE_RSA_WITH_AES_256_CBC_SHA (0x0039)
    Cipher Suite: TLS_ECDHE_ECDSA_WITH_RC4_128_SHA (0xc007)
    Cipher Suite: TLS_ECDHE_RSA_WITH_RC4_128_SHA (0xc011)
    Cipher Suite: TLS_RSA_WITH_AES_128_GCM_SHA256 (0x009c)
    Cipher Suite: TLS_RSA_WITH_AES_256_GCM_SHA384 (0x009d)
    Cipher Suite: TLS_RSA_WITH_AES_128_CBC_SHA (0x002f)
    Cipher Suite: TLS_RSA_WITH_AES_256_CBC_SHA (0x0035)
    Cipher Suite: TLS_RSA_WITH_RC4_128_SHA (0x0005)
    Cipher Suite: TLS_EMPTY_RENEGOTIATION_INFO_SCSV (0x00ff)
```

The server however supports these cipher suites:

```
PORT STATE SERVICE
25252/tcp open unknown
| ssl-enum-ciphers:
| TLSv1.0:
| ciphers:
| TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA (secp256r1) - A
| TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA (secp256r1) - A
| TLS_DHE_RSA_WITH_AES_256_CBC_SHA (dh 2048) - A
| TLS_DHE_RSA_WITH_AES_128_CBC_SHA (dh 2048) - A
| TLS_RSA_WITH_AES_256_CBC_SHA (rsa 2048) - A
| TLS_RSA_WITH_AES_128_CBC_SHA (rsa 2048) - A
| TLS_RSA_WITH_3DES_EDE_CBC_SHA (rsa 2048) - C
| TLS_RSA_WITH_RC4_128_SHA (rsa 2048) - A
| TLS_RSA_WITH_RC4_128_MD5 (rsa 2048) - A
| compressors:
| NULL
| cipher preference: server
| warnings:
| Ciphersuite uses MD5 for message integrity
| TLSv1.1:
| ciphers:
| TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA (secp256r1) - A
| TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA (secp256r1) - A
| TLS_DHE_RSA_WITH_AES_256_CBC_SHA (dh 2048) - A
| TLS_DHE_RSA_WITH_AES_128_CBC_SHA (dh 2048) - A
| TLS_RSA_WITH_AES_256_CBC_SHA (rsa 2048) - A
| TLS_RSA_WITH_AES_128_CBC_SHA (rsa 2048) - A
```

```
| TLS_RSA_WITH_3DES_EDE_CBC_SHA (rsa 2048) - C
| TLS_RSA_WITH_RC4_128_SHA (rsa 2048) - A
| TLS_RSA_WITH_RC4_128_MD5 (rsa 2048) - A
| compressors:
| NULL
| cipher preference: server
| warnings:
| Ciphersuite uses MD5 for message integrity
| Weak cipher RC4 in TLSv1.1 or newer not needed for BEAST
    mitigation
| TLSv1.2:
| ciphers:
| TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 (secp256r1) - A
| TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 (secp256r1) - A
| TLS_DHE_RSA_WITH_AES_256_GCM_SHA384 (dh 2048) - A
| TLS_DHE_RSA_WITH_AES_128_GCM_SHA256 (dh 2048) - A
| TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384 (secp256r1) - A
| TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256 (secp256r1) - A
| TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA (secp256r1) - A
| TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA (secp256r1) - A
| TLS_DHE_RSA_WITH_AES_256_CBC_SHA (dh 2048) - A
| TLS_DHE_RSA_WITH_AES_128_CBC_SHA (dh 2048) - A
| TLS_RSA_WITH_AES_256_GCM_SHA384 (rsa 2048) - A
| TLS_RSA_WITH_AES_128_GCM_SHA256 (rsa 2048) - A
| TLS_RSA_WITH_AES_256_CBC_SHA256 (rsa 2048) - A
| TLS_RSA_WITH_AES_128_CBC_SHA256 (rsa 2048) - A
| TLS_RSA_WITH_AES_256_CBC_SHA (rsa 2048) - A
| TLS_RSA_WITH_AES_128_CBC_SHA (rsa 2048) - A
| TLS_RSA_WITH_3DES_EDE_CBC_SHA (rsa 2048) - C
| TLS_RSA_WITH_RC4_128_SHA (rsa 2048) - A
| TLS_RSA_WITH_RC4_128_MD5 (rsa 2048) - A
| compressors:
| NULL
| cipher preference: server
| warnings:
| Ciphersuite uses MD5 for message integrity
| Weak cipher RC4 in TLSv1.1 or newer not needed for BEAST
    mitigation
|_ least strength: C
```

Neither the server nor the client make use of cipher suites which have been broken in practice yet. They seem sufficiently strong to not be useful target for an attacker.

2.3.7 Webinterface

Detention center employees have access to a webinterface for the pilot system. To access it, they need to log in with secret credentials that are the same for all employees. In this interface, a world map (based on Open-

StreetMap) is visible and can be zoomed and dragged (as shown in figure 3). On this map, pins are visible that correspond to specific “rooms”. A “room” is a user identifier for the purposes of this pilot system, because different detainees can use multiple Nymi ID’s over the duration of their stay in the detention center. Personal names of detainees are not used because of privacy considerations. An employee is aware of which detainee stays in which room of the center, so he also knows the identity of the detainee on the screen.

Around a pin, a circle is drawn to show the accuracy of the measurement, indicating that the detainee is inside that radius at that time. There is also the option to focus on a specific “room” (i.e. active user) on the map using a list of all pins which is always visible. Hovering over a pin shows the battery level of the device (this is also shown in the list of rooms which is always visible) and how long ago the time of last communication was. A green pin indicates that the S & T servers have received a heartbeat from that device within the last 2.5 minutes. If the last communication was more than 2.5 minutes, but less than 5 minutes ago, the pin colour turns orange. If the S & T servers have had no messages from a specific device for more than 5 minutes, the pin colour turns red, a flashing message pops up and an alarm goes off. A grey pin is used to indicate inactive rooms. The webinterface has a fixed amount of rooms it displays, and not all of them have to be in use.

There is also an additional web page in the interface in which the path of a specific selected ”room” can be traced out over time. You can also select the option to show warnings in the playback of the path as well. The website uses javascript for all interaction.

2.3.8 Webinterface - Server Communication

To log in to the web interface, `HTTP` Basic access authentication is used over `HTTPS`. The content of the website itself is also sent over `HTTPS`. The server supports the same available cipher suites for the webinterface as it does for the `TLS` connection in the case of the phone communication (see 2.3.6). Periodically, the web page uses an `AJAX` request (also over `TLS`) to fetch the most recent information with regard to the active devices, which it then reflects in its interface.

2.3.9 Helper application

The helper application was not part of the materials that we were provided with, but based on information in instructional documents and other parts of the pilot system it became clear what it is used for. Using this desktop application, employees of the detention center are able to link a specific Nymi ID to a “room” in the webinterface. The employee knows of this Nymi ID

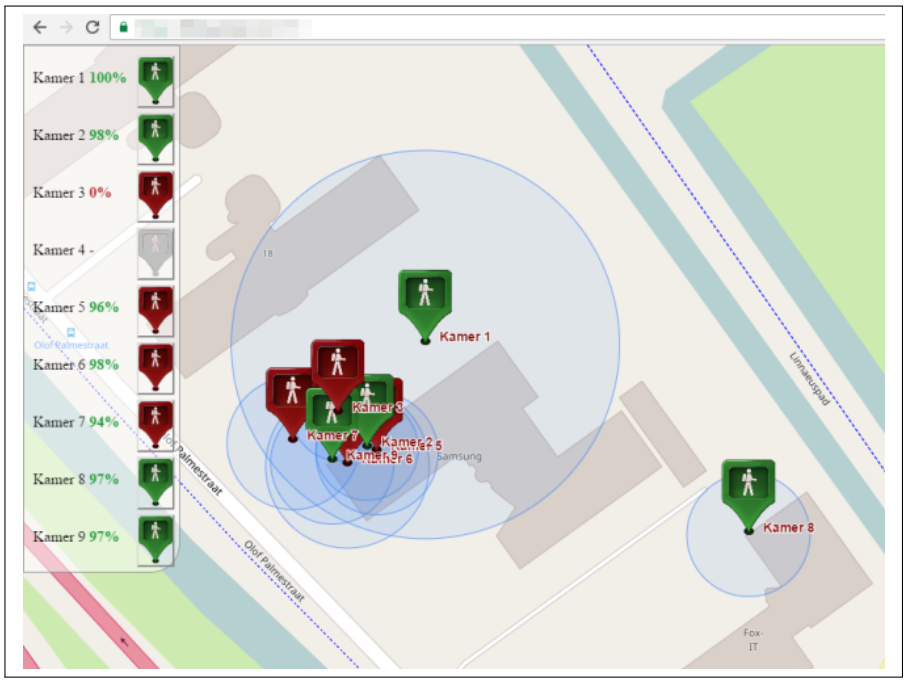

Figure 3: A cropped screenshot of the webinterface displaying mock-up location information.

because it is displayed in the tracking app on the locked Android phone given to the detainee. Before the phone is locked and given to the detainee, the detention center employee matches the ID on screen with the one in the helper application and assigns it to the room number of the detainee. From thereon the webinterface will display the data corresponding to that specific Nymi ID as (location and other) information for the chosen room in its UI. It seems fair to assume the communication between the helper application and the S & T servers is encrypted and happens with `TLS`, just like all other internet traffic that is part of the pilot system.

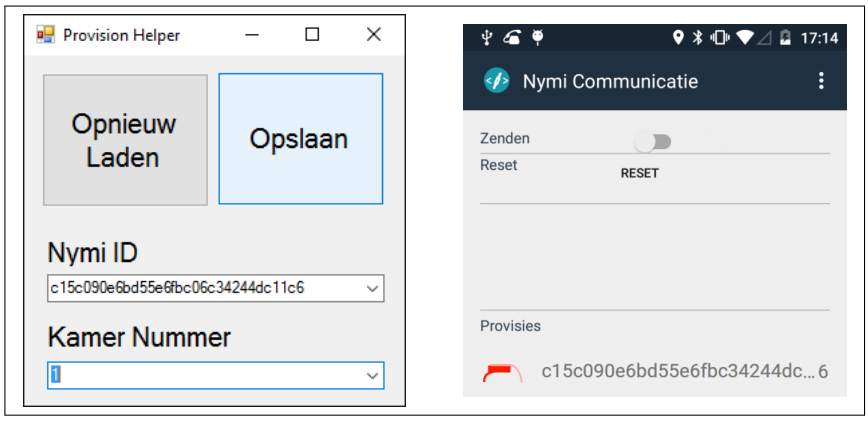

Figure 4: Helper application, running on an employee workstation (left) and tracking app, running on the locked Android phone (right)

2.4 How leave works in youth detention centres

Detainees take part in a process which aims to improve their behavioural and mental situation and facilitate their re-introduction in society. Education (whether it be in the facility or at a conventional school) is a part of this process. Another part can be leave, in which a detainee for instance visits friends or family, practices a sport or hobby, or goes to school or works at a job. The goal and approach of the leave varies between detainees: there is no one-size-fits-all approach and every case is constructed through (and after) discussion between the behavioural scientists at the facility and the detainee in question.

Detainees who take part in a team sport during their leave could do this to build a healthy friend network, to build a feeling of self worth or for instance fill a gap left by addiction. General goals of leave are: the search for a meaningful way to spend the day, to orient oneself on living on ones own, to go to school or to take part in a social network (in the non-digital sense of the word). For every detainee, a risk analysis of the requested leave is made and it must be demonstrated that this specific leave has a valuable role in the treatment of this detainee. Every leave needs to be individually approved by mandate of the director of the detention center.

A detainee starts with one-day accompanied leaves and works towards one-day unaccompanied leave or multiple-day unaccompanied leave in the long run (multiple-day leave is at most one night outside of the facility). For a detainee to start with leaves or progress to a less controlled stage of leave, approval of the Ministry is necessary, based on the dossier and analyses for that detainee. The destination of leave has to be inside of the Netherlands, because they are not allowed to leave the country outside of extremely rare cases with express permission of the Ministry. Depending on

the risks involved in the destination of the leave, the choice can be made to still have it be an accompanied leave.

Whenever possible, the facility tries to have a contact person at the destination of the leave in the form of teachers, sport coaches, family members or managers. They are informed of the situation and are expected to keep an eye on the detainee during leave.

When the detainee goes on leave, they have their picture taken. One of the reasons for this is so this can be used in a police description in case of escape. A detainee on leave is also supposed to keep a "leave pass" on their person at all times. This pass has (among other information) the destination of the leave and the start and end time on it. This pass can be used to justify themselves not being in the detention center if they get stopped by police, because they should show up as being detained when ran through the police system.

After returning from leave, the detainee writes an evaluation together with one of their detention group sociotherapists at the facility and this evaluation is then later discussed with a behavioural scientist. New leaves can only be approved after all earlier evaluations have been discussed.

2.4.1 Escape protocol

If a detainee does not adhere to the agreed upon times of his leave, they are supposed to call in about this in advance with a valid reason. Depending on the situation, the facility may decide to send an employee to their last known location. If a detainee does not call in, they will be called on their personal phone. If they don't answer the phone, the facility will try to reach the other parties involved in the leave (e.g. family members, their coach or teachers) to get information about the whereabouts and situation of the detainee. If the center is reasonably sure that no escape has taken place, because they have for instance had contact with the detainee or contact people, they are may decide to make a notification of a special occurrence if the detainee is more than one hour late. Whether or not this is done is based on the discretion of active management. For each of these decisions, other members of management have a controlling role in deciding if the call was made correctly.

When there exists a reasonable suspicion that the detainee has escaped, or they are otherwise untraceable, investigative services are notified through the GRIP (Gecoordineerde Regionale Incidentbestrijdings Procedure, which stands for Coordinated Regional Incident-management Procedure). Other governmental services handle the escape from here. Depending on the estimated risk attached to the detainee, they will be put on domestic and international wanted & missing persons lists. This means that certain locations in their personal life (i.e. their parents, friends or significant others' houses) can be shadowed, which is the first step that police take. This is of-

ten enough to catch most escapees. If the shadowing of these locations does not lead to results, a more active form of search will be started. Escapees should get arrested during traffic checks or when asked to identify themselves (including at border control, i.e.: before air travel or when leaving the Schengen zone). Interestingly, being detained in a youth detention center is not enough to be put on a list which gets checked at border control.

3. Research

In order to answer our research question,

> *To what attack scenarios is a system such as the one tested at the DJI subject and/or vulnerable and how can these attacks be prevented or have their success impact reduced?*

We will first begin by describing possible attacker models in relation to the pilot system when applied to conventional youth detention centres in Section 3.1. Section 3.2 will propose some practical limitations of the pilot system that could, aside from hindering the utility of the system, be used as a cover of plausible deniability by an attacker. In Section 3.3 and 3.4 we will discuss general vectors for attacks and then focus on analysing specific ways that these could be exploited, and how to remedy those attacks.

3.1 Attacker model

In looking at possible attacker models, we consider the model of a detainee in a youth detention center looking to subvert the system. Instead of looking at the detainees considered to be low risk in the small scale experimental annex where the application of the pilot system is temporary, it is more useful in the long run to investigate attacker models for the pilot system when applied to normal youth detention or general detention. We have chosen to look only at youth detention in this thesis because of time constraints and because the small scale facility is an annex of a youth detention center. A detainee trying to subvert the system could either be working on his own, with help from friends or a criminal organisation. Help from a foreign government is also a theoretical possibility, but seems unlikely in the case of youth detention. If a specific attack has a high up-front monetary or timewise investment that results in something that can be reused (acquired hardware / developed software / acquired key information) this should be considered a lower barrier to entry. This is because criminal organisations could reuse it for other detainees or sell it to other parties. To prevent reuse of attacks it must be possible for detention centres to work without the system as well, so it can be turned off and re-evaluated. One factor that should not be discounted is that effort-intensive attacks might not be worth it if the detainee is going to be released after one more year of good behaviour (unless they are in a hurry). Such attacks are therefore more likely to be considered in the case of open or semi-open penitentiaries.

3.1.1 Goals

To have a clear view of the goals a detainee could have with respect to subversion of the monitoring system, it is useful to enumerate the functions that the system fulfils for the detention center:

- To prevent detainees from not returning to the detention center at all. This is accomplished by knowing whether or not an escape is happening very fast by the system by either detecting a lost signal or an employee keeping an eye on the movement of the detainee. Furthermore, in case of a lost signal, the detention center knows the last known location of the detainee.

- To detect detainee behaviour that does not follow the conditions of leave that were agreed upon (i.e. are factors with a negative influence on their treatment process). Employees of the detention center are able to recognize if a detainee has for instance left their sport training earlier that day / has gone to a casino / taken a route to school which passes a coffee shop / visited family members which are considered to reinforce the problematic behaviour of the detainee. This is accomplished by an employee who has such information about the detainee noticing the location / path of the detainee on the map.

This results in two clear-cut categories of attack results:

- Attempt to escape detention and have the detention center find out as late as possible (at the end time of the leave) instead of within around fifteen minutes. This means that the system does not need to return to a stable state after the attack.

- Return to the detention center after leave, but not follow the agreed upon terms of leave by visiting forbidden locations without the center noticing. This means that the system needs to return to a stable state after the attack, to prevent tampering being noticed after returning.

3.1.2 Scope of attack

We consider two kinds of restrictions that work upon the attacks that affect the monitoring system. These are restrictions in personal movement, communication and movement of goods that are a consequence of a detainee being in a detention center and restrictions concerning the means of attack / attack scope.

1. Restrictions related to the detention center:

    - No access to a cellphone (but sometimes successfully smuggled in).

- Restricted internet access.
- On the outside, detainees have a cellphone and full internet access.
- Incoming mail is checked (for contraband), outgoing mail is not (for privacy). It is fairly easy to think of a way to avoid getting mail checked by falsifying certain addressing information.
- Periodically the halls of detention groups can be scanned with a "mobifinder", i.e. a device that detects active mobile phones.
- Specific sections of the detention center have metal detectors.
- When a detainee exits or enters the detention center he may not set off the metal detectors. When a visitor enters the detention center they may not set off the metal detectors. A visitor can however leave the detention center setting off the metal detector (the guards cannot search them).
- The detention center has airing grounds / gardens (open air enclosed areas) that the detainees can choose to visit each day.
- Sometimes a detainee can be eligible for accompanied leave and have it approved, but still won't be able to go, because of staff shortages at the detention center. (More about this under section 3.2.3)

2. With regards to the means of attack we consider the following restrictions:

    - No restriction to limit of technical knowledge on the attacker side. Some detainees are educated in terms of coding / decompilation / network sniffing and they can receive help from people that can be more educated then them, either by knowing them or hiring them.
    - No time restriction for developing the attack outside of the fact that waiting out your treatment (or in the case of penitentiaries: sentence) is also a very effective way of escaping or being able to do whatever you want.
    - In an attack, there may be no lapse of communication for more than 5 minutes, because that makes the alarms go off, unless the detainee can present a reasonable explanation for this when he is first called by the detention center. (More about this under section 3.2.2)
    - With regards to monetary restrictions: It is hard to quantify this, but it should be taken into consideration that attacks with a high up-front cost may not necessarily be off the table if the attack is reusable or can be sold to other parties (in concept or materials).

3.1.3 Capabilities

We consider an attacker to be able to:

- Have access to all publicly available software, CVE's and services.
- Have access to all publicly available hardware within their budget.
- Have access to a botnet, either be them or an associated organisation owning one, or by renting the service.
- Use a drone to airdrop a sufficiently light object into the detention center airing grounds.
- Perform a bluetooth Man in the Middle attack.
- Perform a WiFi MitM attack.
- Not perform a 3G/4G MitM attack, because 3G and 4G use mutual authentication. A downgrade attack using an IMSI catcher should however be possible.
- place a network hub on the internet cables outside of the detention center of company servers, i.e. perform a MitM on that internet traffic.

3.2 Practical matters

There are some practical matters with respect to the pilot system that should be taken in account if it were implemented at a larger scale than the small scale experimental facility that play into its security.

3.2.1 Cell coverage in the Netherlands

Even though it is the year 2017 at the time of writing this thesis, there are still villages in the Netherlands with near to no cellphone or mobile internet reception[1]. Furthermore, certain buildings that exhibit properties of Faraday cages or the underground areas parking garages will not have reception. Any detention center should be aware of these spots when using the system, because it could lead to false positives with respect to detainees attempting escape, lessening the conclusive value of the alarm sounding. Furthermore, any attack in which it is necessary to temporarily drop from the system to for instance tamper with the locked phone, could be explained away under the excuse of loss of reception.

[1] https://nos.nl/artikel/2140588-tientallen-dorpen-hebben-nog-altijd-nauwelijks-mobiel-bereik.html

3.2.2 Staffing

In youth detention, the situation often occurs that there are more detainees on unaccompanied leave then there is staff available to retrieve these detainees. This is logical, because it does not often occur that a staff member needs to retrieve a detainee on leave, there are a lot more detainees than there are staff members associated to these detainees and it would also not make sense to keep all this personnel on call if they are not needed.

A staff member would for instance travel to the location of a detainee on leave if that detainee or their contact person requests permission from the detention center that that detainee is allowed to deviate from their leave return time (for instance if their music recital is going to finish late). The detention center can decide to send an employee that knows this detainee to supervise and accompany the detainee back afterwards.

In the case a location monitoring system for detainees is implemented, staff might be needed to retrieve detainees for a different reason. The detention center needs to be aware of the possibility that such a system experiences an outage because of an electricity blackout or internet outage (be it a distributed denial of service attack or incidental). If this were to happen, the detention center would need to be able to retrieve detainees for which this is necessary. If it were to happen that there are detainees who were allowed on leave under the understanding that their location was tracked and who might form a risk factor when not tracked, these detainees need to be able to be retrieved based on the discretion of the detention center in case of an outage. In short: a detention center should weigh the risks in using a monitoring system as a way to let detainees go unaccompanied on their leaves earlier than when such a monitoring system were not a factor.

Another matter related to staffing is the situation of control of compliance with leave terms. It is possible for employees who know the story of a certain detainee and the agreed terms of their leave to check the location history of a detainee after the fact to see whether or not these terms were followed. However, it is not so that this informed employee is continually checking the interface, because they can be doing other work or just be done with their working day. The employee who is trusted to watch the interface live, which could be a guard, can hardly know all dossiers by heart. This means that the logging of location (in contrast to the detection of the lack of phone heart beat) is only reliable in fulfilling a controlling function after the fact and does not necessarily provide any actionable information in the present, even though the interface displays live information.

3.3 Software attacks

3.3.1 GPS spoofing

Android gets the location information that it provides through it's location API through the Global Positioning System. GPS consists of a collection of satellites that transmit signals back to earth. Receivers of these signals can work out what their own location is because of the path of the satellites is known, the Doppler effect and math. It is however also possible to spoof these signals. As long as a device receives a collection of signals that also could have been sent by the GPS satellites (and the signals transmitted by the satellites is not signed and deterministic, so these signals can be calculated by other parties) a smartphone won't be able to tell the difference. To make a device see the spoofed signals over the real ones, it is only necessary to make sure that the spoofed signals are stronger than the real ones. For years, dedicated GPS spoofing hardware has existed. The only limiting factor for attackers here would be the price point, which used to be around thousands of dollars, which is still somewhat reasonable if an attacker would be able to get more value out of it by reusing it. In 2015 however, the researchers Huang Lin and Yang Qing presented a cheap working proof of concept for GPS spoofing during the DEF CON 23 conference.[2] Their solution makes use of a bladeRF, which is publicly available for around four hundred dollars. A smart attacker would be able to reverse engineer what the code for this would look like from the presentation slides. However, it is not necessary to be a smart attacker here. Even though the code for their proof of concept was not published as is, it is available through another project that one of the researchers based on their work, which is intended for cheating at the smartphone game Pokemon Go. This code is publicly available. (`https://github.com/JiaoXianjun/poke-move`) This software provides an attacker with exactly the same functionality as the proof of concept presented at DEF CON 23. If an attacker were to carry this hardware (a small laptop and a bladeRF) around outside the detention center they can be somewhere completely different than the tracking app will report to the company servers. A different way of GPS spoofing that is somewhat less interesting in the context of the pilot system is the built in Android functionality for apps to provide a so-called "Mock location" to other apps, which another app will receive when they call the location API. This means that an app which plays back fake locations would also work, but device access is necessary to install such an app. In the case that an attacker has access to the locked phone, many other attacks are possible than just faking the location to the tracking app.

This attack is pretty hard to prevent. It completely takes place outside of the detention center and it is not necessary to take the hardware back inside. One option could be to use the cell tower the phone is connected to

for 4G to falsify certain locations, but it is not clear whether this information is easily attainable through the cell provider. Even if this were possible, it still allows an attacker to go wherever they want within the range of the cell tower.

Something that S & T could look into is the Galileo satellite navigation system that is currently under development. This navigation system has an encrypted version just like GPS, but where the encrypted high resolution version of GPS is only available to the United States military, the encrypted version of Galileo is commercially available.

Result	Effort	Likelihood	Fixable
Delayed noticing of escape attempt or being able to deviate agreed terms of leave and return.	Not very high. The biggest limiting factor here is the purchase of the hardware which is not that expensive. The software is freely available.	Very realistic.	Not easily fixable, but can be limited in effectiveness through falsification of location using the active cellphone tower connection.

3.3.2 Heartbeat spoof

If you recall, in the `TLS` communication of the locked Android phone to the S & T servers only server authentication takes place. We know the communication format that the tracking app on the locked phone uses to pass information to the servers. Also, we know the universal password in this format because it is hardcoded into the app. This means that we can ourselves create messages following this format which the server can not distinguish from other legitimate messages in the pilot system. No IP whitelist is used by the server to decide if a message is legitimate. Adding this would be beneficial, even though that too can be spoofed (though not easily). The only problem that we are faced with here is that we do not know which NymiID to use in our messages so as to pass them for ones from the user that we want to spoof for. If we knew what NymiID corresponded to the one in the app on the locked phone, we could just power the phone down and have our own program take over the heartbeat function with fake location info. In the locked Android phone, the NymiID is obtained through the "Nymi provision" file which is stored under the root user directory.

We will now consider three ways to get information about the NymiID:

(1) No file access or circumvention of the screen lock:
Let us assume that an attacker is not able to get access to any files on the phone or is able to circumvent the screen lock in any way. If a detainee is able to look at the screen of the phone when it is open on the tracking app, they can see the NymiID which is displayed there. This app is open when their supervisor is linking the phone to the band and matching the NymiID with the values in the desktop helper application.

Even if a detainee were able to remember 32 characters by heart, they may not necessarily be able to get all the characters from looking at the screen. Because the app is fixed in portrait mode, has a limited resolution and the font is not monospaced, the amount of visible characters (and how many are cropped) is dependent on the characters that that specific NymiID consists of. Depending on how many characters the detainee is able to obtain it might be able to spoof the heartbeat for the state space of possible NymiIDs, but when the amount of missing characters gets too high ($\geq$ 4 characters) it will become hard to spoof a message for all the possible NymiIDs within the 5 minute alarm window. For 4 missing characters it is already necessary to send $\frac{36^4}{5*60} \approx 5600$ fake heartbeat messages every second, which is very likely to be recognized by S & T. For 3 missing characters this works out to around 155 messages a second, which might still be possible.

Even though this attack is pretty unlikely in the sense that a detainee would have to remember a long string in a very short time, it is also reasonably easy to fix. If the app were made to never try to display the entire NymiID, but drop the last 4 or 5 characters, it would be practically impossible to spoof a heartbeat message correctly based on the displayed information.

Result	Effort	Likelihood	Fixable
Delayed noticing of escape attempt or being able to deviate agreed terms of leave and return	High, because it is very hard to quickly remember a large random string of characters. Also might require a large active message spoofing effort.	Low, but not impossible	Easily fixable.

(2) Circumventing the screen lock and (3) Getting file access
If an attacker were to be capable of getting past the lock screen, the security of the pilot system would be compromised. Not only would they be able to install an app to fake GPS location to the tracking app, they would also be capable of enabling USB debugging, getting the tracking app from

the phone and installing their own edited version based on the decompiled tracking app to take over the function of the normal tracking app on their next leave (when a new provision file is generated for the new app). The icon for the real running tracking app could be hidden and replaced by the icon of the new fake app so the supervisor selects that one the next time instead. This modified app could itself lie to the server, or just pass the NymID on to the attacker.

This is all well and good, but how do you get past the screen lock if you do not know the password? Either by using software to remove it or bypassing the lock screen in it's entirety by getting access to all (or a part of) the system memory (and thus the provision file). Whether any of this is possible depends on the model of phone and Android version running on it. For the combination of model and version that we were provided with (ZFD VFD 600 running Android 6.0.1 at May 1st 2016 patch level) there are no ready-made tools or recovery images (like TWRP[2] or ClockworkMod[3]) to flash to the device which would allow an attacker to retrieve the provision file or dump the memory of the tracking app process. For a lot of other smartphones this is in fact the case.

To make sure this is not possible in the future, great care should be taken to research and monitor whether software or recovery images are (/become) available.

Furthermore, the Android version running on the phone should be kept updated to the most recent version as much as possible, because CVE's for Android regularly come out[4] which could provide enough access to the device to break the system.

Saving the provision file under the Android Keystore does not improve the security by much, because the NymiID is still part of the process memory when the app is running. It would however protect against attacks that are only able to get files from the phone and not perform memory dumps.

Result	Effort	Likelihood	Fixable
Delayed noticing of escape attempt or being able to deviate agreed terms of leave and return	Low in the case of publicly available software and recovery images. Extremely high in the case of CVE's without published exploits.	Very low, if these attack solutions are well-monitored.	Not needed nor possible, but requires frequent updates and monitoring of attack solutions.

[2] https://twrp.me/

[3] https://www.clockworkmod.com/

[4] http://www.cvedetails.com/vulnerability-list/vendor_id-1224/product_id-19997/year-2017/Google-Android.html

3.3.3 DDOS

A distributed denial of service attack, in which many sources flood the S & T company servers with communication at the same time, is a way to prevent normal network usage in the pilot system. This way, near to no heartbeat messages will make it through and all users of the system will drop at the same time. The only difference between this and a detainee cutting removing their Nymi band or throwing away the locked phone is that the detention center cannot determine which of the detainees that are part of the pilot system are attempting escape, and which are not, but just dropped from the system. It is still directly perceivable that something is happening, but it requires a lot of effort on the side of the detention center to keep track of all the detainees that are now not being monitored by the pilot system any longer by having staff members making calls or retrieving detainees. Depending on the way a DDOS is handled, an attacker may have more or less time before their actual escape attempt is recognized. If a DDOS attack were to last days, which is possible, the entire system would be unavailable for that entire duration. An attacker with access to a botnet or a booter / stresser service (a service offering a botnet for hire) could easily perform a DDOS attack. Using a botnet as a service can be done relatively cheaply[5]

A DDOS attack is not really preventable. The ISP may have DDOS prevention or detection mechanisms in place to help. The detention center should still prepare a procedure to follow in case of a (pilot) system outage.

Result	Effort	Likelihood	Fixable
System outage, which could be used as a cover for an escape attempt. It can also be used to not follow agreed terms of leave, but it seems unlikely to be used this way.	High, unless the attacker already has access to a botnet.	Could happen, but unlikely.	Not really fixable, but should be prepared for.

3.3.4 Active Bluetooth MitM

The trust in the communication between the Nymi band and the locked phone is based on the keys that are exchanged at the start of their pairing using the Diffie-Hellman key exchange. One of the properties of the DH key exchange is that it is not resistant against an active Man in the Middle. If

[5] `https://blog.radware.com/security/2017/03/cost-of-ddos-attack-darknet/`

there is a party that sits between both communicating parties who performs a DH key exchange with both sides, they will not see the difference, but the MitM will be able to read and change communication.

In the case of the pilot system this would mean that a device like a raspberry Pi would be able to see the messages between the Nymi band and the locked phone. Whether or not this device would be able to replay messages, generate them or change them depends on the choices made by Nymi in their proprietary protocol (which is not publicly available).

In the Nymi whitepaper[1] the following is said about active MitMs between the Nymi band and Nymi Enabled Applications:

The protocol also guarantees that any active 'man-in-the-middle' (MIM) adversary must continue to actively modify radio communication between the Nymi Band and the third party device at all times to avoid detection and termination of the pairing. Active MIM attacks in a radio communications setting are extremely difficult and costly to perform.

In reality, active MitM attacks on Bluetooth are not extremely difficult or costly to perform. An easy way to do this would be to use a raspberry Pi running Raspbian (40 euro or less) and publicly available software for Bluetooth MitM[6]. The only obstacle for an attacker to perform a MitM here would be to reverse engineer the protocol that Nymi uses specifically and to find out if it is possible to replay or generate their own messages. If this is possible, an active Bluetooth MitM attack on the pilot system would be possible in theory.

If a detainee were to be able to have such a device on them when they get fitted with the Nymi band and the band is linked to the locked phone, they would be able to perform this attack and have their device always report an authenticated status, even if the Nymi band loses this status. Getting the locked phone to keep connected to a MitM device from inside the detention center to the supervisors office and outside is extremely hard. We will discuss how here because it is very involved and challenging, but not because it is in any way likely to be able to be executed in practice.

Getting this device inside would have to happen either through a bribed guard, or through a drone air drop to the airing grounds. Then the detainee needs to be able to carry this device from the airing grounds to the supervisors office without passing through a metal detector. If this is possible depends on the layout of the detention centre in question. Furthermore, depending on whether or not stationary or hand-held metal detectors are used at the exit, another method would need to be used to get the device out of the centre. In the case of a stationary metal detector, the detainee would be expected to set off the system, so they can just walk out with the

[6] `https://github.com/conorpp/btproxy`

device. If hand-held metal detectors are used, things get harder still. In that case, if it is possible based on the layout of the building and the exit metal detectors, the attacker needs to leave the first MitM device inside to not have it detected by the manual metal detection, while at the same time staying in range of a similar device on the other side of the metal detectors that is passed the correct key information by the first device, so that second device can take over the MitM function. If all this has been successful, the detainee can remove their Nymi band and give the phone and MitM device to an accomplice who will take over their routine.

Whether or not a detainee would be able to keep the MitM going when returning to the detention center after use on leave depends on the procedure in place for returning the Nymi band and locked phone, but it seems extremely unlikely. This procedure is not described in the instructional documents we received. It is logical to assume that a detainee would not be able to set off stationary metal detectors when returning, because they would be able to smuggle anything back in. More logical in this case would be that they have to drop off the Nymi band and locked phone somewhere at the entrance of the facility, before passing through the metal detectors. In this case they would be able to leave the device outside and have their tampering with the system undetected. The same goes for hand-held metal detectors.

This attack cannot be prevented by S & T, but it is probably unnecessary to protect against this attack because it is practically impossible to perform. It could be prevented by placing metal detectors inside the centre on the path between the airing grounds and the group supervisors office.

Result	Effort	Likelihood	Fixable
Delayed noticing of escape attempt or being able to deviate agreed terms of leave and return (once)	Very high	Very low, but theoretically possible.	Not needed.

3.3.5 Bad certificate Webinterface MitM

The tracking app (and likely also the helper application) throw exceptions when presented with an invalid certificate, preventing the connection from being established. This is not the case for web browsers. These give a warning prompt to the user that the connection is untrusted, but users are sometimes inclined to simply ignore such warnings. If the attacker is able to establish a MitM between the employees workstation and the company servers ánd this employee ignores the warning, the attacker has full control over the locations displayed on the screen of the webinterface. A MitM of this kind could either consist of a tap on a network cable outside of

the detention center or company servers, or a rogue device on the inside of either the detention center network or company network using ARP flooding. This device could either be a hacked computer (possibly even the pc on which the webinterface is requested itself) using for instance a malicious usb or email attachment, or a device like a raspberry Pi which has been airdropped to the airing grounds of the detention center which has been attached to an unprotected Ethernet port. All the MitM attacks inside of the detention center or company network would only work if the network layout is beneficial and there are no protective measures for internal communication preventing ARP flooding.

This attack is preventable through user education and by adding a warning to the instruction manual, but it is not fully preventable this way. A clear link should be made between the browser warning and a possible hacking attempt. Even though it might seem a solution to sign all location updates with the current time appended (to prevent replay) and verifying these messages client-side, this would not work, because the webinterface itself is also data that is transmitted to the client. This means that the attacker can just modify the check to always succeed. A solution that would work is to have an application instead of a website which would prevent usage in the case of an invalid certificate, but this is a lot less versatile in terms of usage.

Result	Effort	Likelihood	Fixable
Delayed noticing of escape attempt or being able to deviate agreed terms of leave and return	Getting an external MitM up is a lot of effort, but so is an internal one. The software for this can be written in a couple of days and the monetary cost for hardware is lower than one hundred euro.	Realistic.	Easily fixable.

3.4 Hardware attacks

3.4.1 Stolen NCA phone

An NCA phone is the only part of the pilot system that is capable of bringing a Nymi band into the authenticated state. If a detainee were to acquire a phone with the NCA app on it (the NCA app is publicly available, because it is part of the Nymi standard software) and the key/biometric information for the profile on the Nymi band, they would be able to reauthenticate the Nymi band using heart rate. Lacking a way to steal the key information from the NCA phone in the detention center, stealing the corresponding NCA phone itself will have to do. In the instruction manual for daily use there is nothing about safe storage of the NCA phone. It only says that it should be put back on the charger after authentication has been successful. Unless part of the instruction manual, it cannot be conclusively said whether this charging happens in a safe compartment. Taking two phones outside instead of two could be tricky, but would maybe be possible. When a detainee is outside and is in possession of the NCA phone for his band, he can just remove his band, turn it inside out (which is possible because the band is very flexible) and then reauthenticate using heart rate with his arm on the outside of the band instead of inside. This might be hard to accomplish within five minutes, because the heart rate detection of the Nymi is sometimes inconsistent and very much influenced by mental state and body position. After a successful reauthentication, the Nymi band can be given to an accomplice who will take over the route that was to be part of the leave of the detainee to prevent detention center employees from noticing anything is different. It would maybe, but unlikely be possible to also return to the detention center without anything being noticed as being amiss (put the band back on, reauthenticate), but it seems unlikely that it would be possible to get both phones back in and also be able to put back the NCA phone where it came from without it either having been noticed earlier during their leave, or when they try to put it back.

This attack is preventable by adding some text to the instruction manual about keeping the NCA phones safe and maybe adding a lockable drawer to the standard procedure. It is likely that the average employee would by themselves already be aware of the risk of the phone getting stolen, but adding it to the procedure seems sensible. If possible, heart rate authentication should also be disabled or at least configured to someone else's heart rate profile. Making sure the NCA phone has a password is also useful, because it requires an attacker to steal the phone when it has just been used by the authenticating party (meaning it has to be stolen under their nose) and also requires the attacker to keep the phone from locking itself because of inactivity.

Result	Effort	Likelihood	Fixable
Delayed noticing of escape attempt (easier) or being able to deviate agreed leave terms and return (harder)	No money or technical knowledge is required. The only thing needed is luck and inattentiveness on the side of staff.	Could realistically happen.	Easily fixable.

3.4.2 Maintaining the continuity circuit during removal

Just like with an ankle monitor, the Nymi band can be cut to make an escape attempt and it will be noticed straight away. This does happen in practice.[7] More interesting is to try and cut the band without the Nymi band itself noticing.

In the official Nymi whitepaper, the company behind the Nymi band makes the following claim [1]:

> *If the Nymi Band is removed from the wrist or cut, the pilot system will immediately detect the intrusion and the Nymi Band will go into the inactivated* [sic] *, preventing access to any internal data.*

Upon dissection of the device the wiring of the continuity circuit in the band shows to be laid out like this:

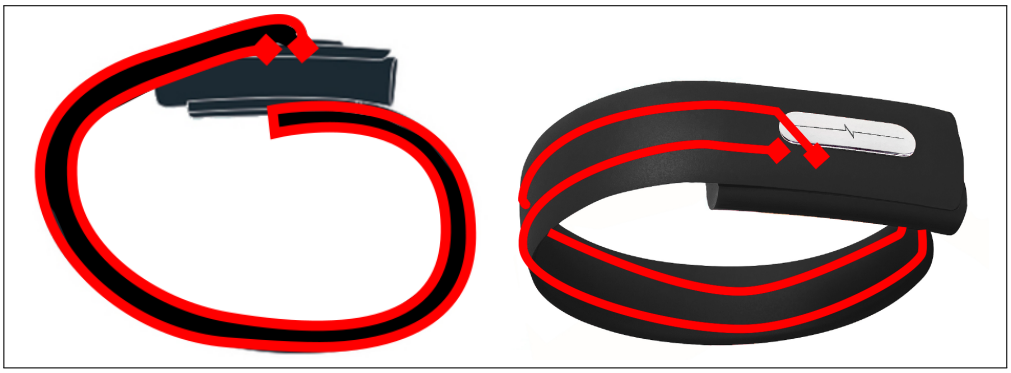

Figure 5: Side view (left) and perspective view (right) of the Nymi band continuity circuit. Wires in red, contact points as diamonds. For the side projection, the wires that are side by side in the real device are shown above one another for clarity.

[7] `https://nos.nl/artikel/2172525-zeker-35-criminelen-knipten-vorig-jaar-enkelband-door.html`

As you can see the continuity circuit is not a closed loop around the arm; Both contact points are on the same end of the band. The flex circuit goes all the way to the other end, up to the point where one can readjust the strap on the underside of the device, and loops all the way back to the top of the device. The reason why you cannot just remove the band at the readjustment point is because of a physical hook preventing the strap from moving to far, not because there is actual wiring there. This means that there is a part of the band where you can cut the strap without cutting the circuit. Even if the strap were to be a loop with a contact point on the underside of the device, the material of the strap can easily be cut open with a utility knife. After this you can melt off the insulating film around the flex circuit wires and solder a kind of "electrical bridge" to two points and cut the circuit between these points. This way the arm can also be removed from the theoretical improved device. We have demonstrated that both these attacks work in practice.

This problem is not trivial to solve. A possible improvement would be to coat all Nymi bands in a hard composite material to make it harder to cut through safely. This would mean that the flexibility of the bands would be lost (not re-adjustable to fit many size any more). An attacker with the right tools could still be able to get past it, but it would at least stop anyone being able to escape using just a utility knife.

Result	Effort	Likelihood	Fixable
Delayed noticing of escape attempt.	Very low.	High.	Somewhat, but not fully.

3.5 Attack overview

Attack type	Result	Effort	Likelihood	Fixable
GPS spoofing	Delayed noticing of escape attempt or being able to deviate agreed terms of leave and return.	Not very high. The biggest limiting factor here is the purchase of the hardware which is not that expensive. The software is freely available.	Very realistic.	Not easily fixable, but can be limited in effectiveness through falsification of location using the active cellphone tower connection.
Heartbeat spoof (no locked phone access)	Delayed noticing of escape attempt or being able to deviate agreed terms of leave and return	High, because it is very hard to quickly remember a large random string of characters. Also might require a large active message spoofing effort.	Low, but not impossible	Easily fixable.
Heartbeat spoof (locked phone access)	Delayed noticing of escape attempt or being able to deviate agreed terms of leave and return	Low in the case of publicly available software and recovery images. Extremely high in the case of CVE's without published exploits.	Very low, if these attack solutions are well-monitored.	Not needed nor possible, but requires frequent updates and monitoring of attack solutions.

DDOS	System outage, which could be used as a cover for an escape attempt. It can also be used to not follow agreed terms of leave, but it seems unlikely to be used this way.	High, unless the attacker already has access to a botnet.	Could happen, but unlikely.	Not really fixable, but should be prepared for.
Active Bluetooth MitM	Delayed noticing of escape attempt or being able to deviate agreed terms of leave and return (once)	Very high	Very low, but theoretically possible.	Not needed.
Bad certificate Webinterface MitM	Delayed noticing of escape attempt or being able to deviate agreed terms of leave and return	Getting an external MitM up is a lot of effort, but so is an internal one. The software for this can be written in a couple of days and the monetary cost for hardware is lower than one hundred euro.	Realistic.	Easily fixable.
Stolen NCA Phone	Delayed noticing of escape attempt (easier) or being able to deviate agreed leave terms and return (harder)	No money or technical knowledge is required. The only thing needed is luck and inattentiveness on the side of staff.	Could realistically happen.	Easily fixable.
Maintaining the continuity circuit during removal	Delayed noticing of escape attempt.	Very low.	High.	Somewhat, but not fully.

4. Future work

Some points that might be interesting to look at in further research:

- The protocol that the Nymi band uses to communicate with Nymi Enabled Applications is not publicly documented. There is documentation for public-facing API methods in their SDK, but the background workings are not explained. We recommend future researchers to try and reverse engineer this protocol and Nymi firmware to find out how secure it really is and what (if any) features it offers on top of confidentiality. Furthermore, the source of randomness in the Nymi band is also very much something to look into.
- With respect to the pilot system: We were not provided with the desktop helper application, so to get a better view of the security of the entire project it would be beneficial to also investigate this part of it. It is also an option that the functionality of this application is integrated with the webinterface, in which case this would no longer be necessary.

5. Conclusions

We have found a collection of attacks that work against the pilot system in the context of youth detention centres, of which some could realistically occur in practice. For these attacks, we have proposed ways to either fully prevent them or limit their effectiveness. We have addressed points on which procedure should be clarified or added to to ensure secure usage of the pilot system.

There are certain attacks that cannot fully be prevented. These are: GPS spoofing and maintaining the continuity circuit while cutting the band.

Our recommendation with respect to the implementation of the pilot system in conventional youth detention centres is that it should be possible if specific parts of the system are improved, and the risks that will still exist are taken account during the decision of whether or not to use the system at what point in the treatment process of a specific detainee.

Bibliography

[1] Nymi Inc. Nymi band white paper, 2015. https://nymi.com/sites/default/files/Nymi%20Whitepaper.pdf.

[2] HUANG Lin and YANG Qing. Gps spoofing, low-cost gps simulator, 2015. https://github.com/op7ic/defcon-23-slides-only/raw/master/DEF%20CON%2023%20presentations/Speaker%20%26%20Workshop%20Materials/Lin%20Huang%20%26%20Qing%20Yang/DEFCON-23-Lin-Huang-Qing-Yang-GPS-Spoofing.pdf.