Counter Expertise Review of the TNO Security Analysis of the Dutch OV-Chipkaart

Introduction

In February 2008 TNO produced a public report summarising its security analysis of the Dutch OV-Chipkaart. The TNO analysis had been commissioned by Translink Systems (TLS) to assess a technical presentation at a recent conference of the Chaos Computer Club (CCC), which claimed that the Mifare Classic card technology and algorithm had been broken and were essentially unfit to underpin the security of the TLS OV-Chipkaart system. Any such claims are damaging to public confidence and when substantiated, would require a significant and costly programme of remedial work for the transport companies and potentially cause distress and inconvenience for travellers. Therefore the TNO report's findings, conclusions and recommendations are extremely sensitive, not only from a business viewpoint, but also for government planning and for public acceptance of the system. For this reason the Dutch Ministry of Transport, Public Works and Water Management authorised an independent Counter Expertise (CE) task to review the TNO report. The selected Counter Expertise Bureau (CEB) was the Information Security Group (ISG) at Royal Holloway University of London and the project was led by the ISG Smart Card Centre (SCC). The CEB was selected on the basis of independence, the information security and smart card expertise of the project team, as well as extensive experience of transport system security as a whole. The Het Expertise Centrum (HEC) was responsible for general facilitation of the task and acted as the interface between the Ministry and the CEB.

The Counter Expertise Task and Adopted Methodology

The CEB was asked to address a number of specific questions in its review and these are summarised below:

- Has the TNO investigation been set up/conducted according to an adequate methodology?
- Is the TNO report complete?
- Are the findings, conclusions and propositions well-founded, correct and univocal?

Particular emphasis was directed towards the TNO report recommendations and whether these would be effective, realistic, practical, complete and future-proof. The CEB was explicitly asked to formulate its own independent professional opinion concerning the findings, conclusions and recommendations in the report. The methodology for the CE task was essentially as follows:

- An initial review of the input set of documents.
- Interviews with TLS and TNO to clarify the contents of the documents and share opinions and experiences relevant to the OV-Chipkaart situation.
- Detailed reviews including cryptographic aspects, attack scenarios and methodology.
- Systematic review of the TNO report against the ministry's requirements.
- Formulation and reporting of the CEB's opinions and recommendations.

The input set of documents included the full TNO report number 34642 (as well as the public excerpt). Consideration was also given to recent publications on OV-Chipkaart attack methods, unpublished research work[1] and other relevant information in the public domain. This was supplemented by confidential documentation from TLS describing the overall system design plus existing detective and preventive countermeasures to combat fraud and security exploits. TNO also provided a memorandum on its follow-on review work for TLS. The output from the CE task is this public report; however all detailed working notes based on sensitive and confidential information have been shared with both TLS and TNO to assist with their further investigations. It should be noted that both TNO and TLS gave full co-operation to the CE process, providing all necessary documents, ensuring appropriate experts were available for interviews and answering all questions that were raised. The CEB would like to thank TLS and TNO for their co-operation in this matter and also the various researchers who offered their information, comments and suggestions.

The TNO Methodology, Scope and Approach

The general methodology adopted by TNO as evident from the detailed report, first investigated the Mifare card technology (that underpins the OV-Chipkaart) and the CCC presentation, before considering the technical consequences of the attacks and then applying these to real-world attack scenarios for the transport system. There were also sections in the report covering risk assessment and the operational management of incidents using back office detection and correction measures. Finally, conclusions were drawn from the work and recommendations were made.

The CEB considers that, the general methodology used by TNO was appropriate for the task and that given the short duration of the task, the report is evidence of a remarkable amount of professional and systematic work. It was correct to consider the problem from a complete system perspective and to consider real-world scenarios, the motivation and execution of criminal attacks, the data and system level security measures as well as the card cryptographic algorithm. The methodology used for risk assessment might have been improved by more closely linking the impact of attacks to real financial value. This is because the financial value may change during the roll-out of the system (and generally over time) and this has an impact on the attractiveness of the system as an attack target and the planning for attack mitigation measures.

The scope of the TNO report is clearly focussed on the impact of the CCC presentation on the continued use of Mifare-based solutions using the CRYPTO1 algorithm and the overall OV-Chipkaart system. It specifically excludes the Mifare Ultralite card that is used for one-off day tickets. From interview discussions, TNO are quite certain that the Ultralite card was not within their task remit as it does not use CRYPTO1. Furthermore, TNO had carried out an earlier investigation for TLS that addressed the Ultralite card, characterised the problems and proposed countermeasures for some exploitation scenarios.

The general approach taken by TNO is regarded by the CEB as “evidence-based”, in contrast to the recent research publications that tend to be more visionary. Whilst the TNO report describes ways that the attacks and the general situation could develop in future, it places more emphasis on what was known, proven, justified and measured at the time of writing. This type of approach is valid in many cases, however because of the fast moving nature of the situation there is the expectation of a flood of follow-on publications showing faster and optimised attacks. Although the TNO report acknowledges this, the CEB would have attached more weight to these factors when making future recommendations.

[1] This included a draft paper from Nohl et al and confidential information from Radboud University Nijmegen

Completeness of the TNO Report

As previously mentioned the general methodology used to generate the report was appropriate and thus the expected classes of content were present within the TNO report. The CEB was specifically asked to comment on whether TNO had given equal attention to all levels of the system:

- Level 0 = The card
- Level 1 = Terminal/reader equipment
- Level 2 = Train stations, bus depot systems
- Level 3 = Public Transport Operator (PTO) systems
- Level 4 = Central Processing Systems

The TNO report covered levels 0 and 4 in detail and level 1 to a lesser, but adequate extent. Levels 2 and 3 were not considered. From interviews with both TLS and TNO it was explained that the missing levels currently have no functional role that is relevant to security and fraud detection. It was however noted that the PTOs receive a duplicate log of travel transaction records, which is currently only used as a double check on the financial clearing process. Because of these duplicate transaction logs and the PTOs' detailed knowledge of their deployed infrastructure, there could be a case for involving PTOs in identifying the locations (e.g. particular stations) where frauds occur.

The description of the Mifare CRYPTO1 algorithm was brief, being based only on publicly available documents, as further detailed information had not been released to TNO by NXP. In the opinion of the CEB, this should not significantly undermine the value of the TNO report, as the most obvious and fundamental limitations of CRYPTO1 have long been well-known in the expert community:

- That the algorithm uses a very short (48-bit) key.
- That the security of the card strongly relies on the secrecy of the algorithm.

TNO and the CEB share a common opinion that reliance on the secrecy of the algorithm is in complete contrast to recommended best practice for cryptographic systems. It is derided in the expert community as "security by obscurity" and is one of the reasons why CRYPTO1 is proving such a popular target for attack. As a result of the latest publications, there is already significant information in the public domain concerning the algorithm, with the prospect of more to come and so additional information from NXP does not seem to be essential. However, if the full algorithm details had been released to TNO, it would have suited TNO's evidence-based approach and reduced the need to speculate on potential structural weaknesses and associated attack scenarios.

The review of the CCC presentation was adequate and this led to more detailed content on the technical and security consequences leading to various attack strategies. There was a considerable amount of content in this section and many techniques were considered. It was not fully complete as at least one attack strategy was omitted and some techniques that could make other attacks much more practical with limited hardware were missing. From interviews it was discovered that TNO have since carried out further analysis on some additional attacks/techniques as part of a follow-on study for TLS and the resulting memorandum was shared with the CEB. The exploitation section, where attack techniques were applied to real-world scenarios, was quite detailed albeit without the missing attack scenarios/techniques. The criminal use of various types of equipment and clone cards was considered and one example business case was presented and discussed. The impact of card attacks on the various stakeholders was also considered, including a brief section on customer privacy.

There was considerable attention to the operational management of security and fraud incidents that required an examination of the many existing validations and report processing rules carried out by the TLS back-end systems. The available mechanisms for disabling invalid or suspicious cards were also considered. The section, whilst very informative, cannot be considered as fully complete because the time taken to respond to individual events was not precisely defined, which is necessary to accurately determine the window of opportunity for a cloned or fraudulently modified card. From interviews with TNO and TLS it was possible to validate the default timing response given in the TNO report, although this was based on the current lightly loaded scenario with very few incidents to deal with. However, TNO and the CEB agree that there is scope for TLS to adapt its staffing levels, automated systems and speed of response, should incident levels rise significantly; but at present there are insufficient statistics to help predict or characterise this rise.

As stated previously, the risk assessment is not considered fully complete as the impacts were not directly linked to financial value and the frequency of successful attacks. Furthermore, the type and technical/financial capabilities of the attacker(s) were not explicitly stated. From interviews with TNO it was confirmed that the probability of incident was based on criminal attack rather than proof-of-concept attacks by researchers. A number of carefully reasoned enhancements to the current system solution were presented to mitigate the risks, although the practicality and effectiveness of these measures have not yet been fully proven.

The Mapping between Findings and Conclusions/Recommendations

Generally the investigation of detailed issues and attacks within the TNO report was carried out in a logical and well-grounded manner. The use of the findings to generate conclusions and recommendations was not solely founded on available evidence, but also relied on assumptions from TNO's past experience.

There is a clear conclusion and a consistent message that runs throughout the report: that the Mifare Classic card will need to be replaced. Some high level suggestions are given on selecting a successor for migration, although the recommendations make no direct linkage to international standards, security evaluation or cryptographic best practices.

Whilst the core conclusion was correct, some assumptions concerning technological barriers to attacks were undermined by the missing attack scenarios and techniques mentioned earlier. These assumptions are critical to estimating the speed at which an attack may be conducted and the associated cost and expertise requirements. The attractiveness of the system as an attack target was not strongly linked to the value of attacks or the deployment phases of the system and this coupled with the technical assumptions put a large tolerance on any time planning recommendations. This tolerance grows even larger when considering that the plans rely on the implementation of additional (currently unproven) remedial measures. Considering the wide tolerance, it is unlikely that the TNO recommendation for the migration readiness milestone is a fully precise estimate for the completion of the complex sequence of activities.

The conclusions and recommendations say little of how public opinion, damage to business reputations and changes to customer behaviour can influence migration planning. Furthermore there seems to be no explicit recommendations that are aimed at “future-proofing” the migrated system.

Review of the TNO Report: The CEB's Conclusions and Recommendations

The CEB concurs with TNO that the Mifare Classic 4k used in the OV-Chipkaart will need to be replaced. The CEB would go further and recommend that any replacement should be based on an algorithm that has been rigorously assessed by the cryptographic expert community, that does not rely on secrecy of the algorithm for security and uses a key length in accordance with cryptographic key length recommendations[2]. As the basis for this, the CEB is convinced that CRYPTO1 has been reverse engineered, to such an extent that it can no longer be regarded as secret. For some time, there have been reports of unlicensed Mifare products in the market that use the CRYPTO1 algorithm. Therefore, with high probability the algorithm was reverse engineered before the CCC presentation. However even if that was not the case, the CCC presentation and subsequent publications suggest that complete disclosure of all final details of the algorithm is imminent. It should not be assumed that publication is solely dependent on the group behind the CCC presentation, as there are other expert teams capable of deducing the final details, now that they have been drawn to the task by the high-profile media coverage. The CEB agrees with statements in the TNO report, that once CRYPTO1 is fully published, it is well within the range of key-cracking equipment that can be used to recover secret keys.

The CEB believes that fast key-cracking equipment is relatively affordable and easy to obtain, but considers that other techniques (such as pre-computation) could provide cheaper, faster and simpler alternatives. Furthermore, the CEB, with the benefit of publications that were not available at the time of the TNO report, considers that the weaknesses being exposed in the algorithm will be exploited to extract keys with significantly less specialised and lower cost equipment. Indeed, the security of proprietary stream ciphers has a reputation for "falling apart" once exposed to scrutiny by the cryptographic expert community. The TNO investigation mainly focussed on an attack facility being developed privately by an individual group with corresponding implications for the required effort, difficulty and time. However, given the high-profile media coverage of the OV-Chipkaart system it is quite possible that multi-skilled collaborative teams will form, with the potential for fast development and optimisation of attacks.

The CEB shares the TNO view that once a card's secret keys are exposed an attacker may modify the card contents or create a copy of the card in an electronic card emulator or in another type of card (clone platform). Extraction of the secret keys presents the maximum opportunity and flexibility for the attacker although in some cases modification of an issued card's contents may prove possible without knowledge of the keys. In common with TNO, the CEB has not been able to find a scenario where knowledge of CRYPTO1 would lead to exposure of a system master key.

The CEB disagrees with the TNO public statement that the cloned cards would necessarily have a significantly different appearance to legitimate OV-Chipkaarts, as having seen a number of real examples of OV-Chipkaart artwork, we believe that some would be relatively easy to reproduce. During the TNO interviews it was clarified that by cloned card in this context, it had meant the hardware emulators that can appear as card-size printed circuit boards, as actual cards were thought less likely to appear. Whilst these hardware emulators may be used in the short term, the CEB considers that one should be well prepared for suitable smart card chips becoming available in the market, which attackers could then embed into convincing looking OV-Chipkaarts. The CEB

[2] Cryptographic Key Length Recommendations from various expert bodies can be found at http://www.keylength.com/en/3/

understands that some of the proposed attack countermeasures may lead to increased emphasis on physical card verification by ticket inspectors, in which case the card body is of greater importance. Currently there seem to be many card artwork designs (and with the prospect of more in the future) and so it may prove difficult for a ticket inspector to visually differentiate between legitimate and counterfeit cards. It is therefore recommended that future OV-Chipkaarts (regardless of which cryptographic algorithm they use) incorporate at least one common and recognisable anti-counterfeit measure, such as laser engraving or a hologram.

TNO concluded that the scenarios for exploiting the system have limited impact on cardholders. The CEB would agree that financial losses would be limited, but impact can also be anxiety resulting from the publication of proof-of-concept attacks, rather than criminal activity. A criminal attack may target stored purse value or its top-ups and whilst the cardholder should be reimbursed for any loss, the experience may again be a source of distress and annoyance. However, to put this into context, a citizen may be at far greater risk of fraud and exploitation when accessing general services via the Internet or telephone.

The CEB would agree with TNO that data privacy is not the major issue for the OV-Chipkaart, as the chip itself currently only contains the cardholder birth date and some travel history, making it a very poor source of personal data compared to Phishing[3] for example. However the CEB appreciates that privacy is a matter of personal opinion and so any perceived threat that is not properly addressed could lead to further loss of confidence.

Current and future attacks might cause the public to lose confidence in the OV-Chipkaart initiative and indeed in other card systems, regardless of the underlying technologies. There may also be negative changes in the behaviour of some customers and the media coverage of proof-of-concept attacks may provide encouragement to try other types of fraud. Furthermore, if the attack techniques become sufficiently simple, some cardholders may be tempted to become amateur attackers to reduce their personal travel costs.

TNO considered losses that might be suffered by a transport company, but put less emphasis on a company's reputation and secondary effects. The impact on the transport companies is not simply measured by the cost of any fraudulent losses, but a range of factors including the effect on business reputation. Unfortunately, as soon as proof-of-concept attacks are publicised, there is likely to be considerable impact that may be partly independent of the level of criminal attacks. From experience with another cloning situation it is known that this can lead to a surge in customer complaints, claims and enquiries that consume considerable customer care and technical investigator resources. It is therefore important to consider that a technology migration decision may be influenced by damage to a company's reputation and added resource costs as well the financial impact of the attacks themselves.

TNO stated in its recommendations, that the “the need for migration is not acute”, but then stressed the need to reach “migration readiness”. During its initial review of the TNO public report, the CEB interpreted these statements as implying a somewhat relaxed approach to migration. However from interviews it became clear that this was not the impression that TNO had intended. For clarity, the explanation from the TNO interview is presented. The TNO reasoning behind the first statement is that there is currently no evidence at all of any such attacks on the TLS system. Furthermore, there are only two locally deployed transport systems using low fare and purse values and so the system

[3] Phishing is a widespread criminal technique using emails to trick users to reveal personal and financial information

does not present a very attractive target and business case for criminal exploitation. There is also an assumption that the remedial/countermeasures proposed by TNO prove practical and are implemented within the system. The "migration readiness" was clarified as meaning the stage when all the PTOs would have deployed the new reader/terminal equipment in all train stations and buses, and so the subsequent migration would be primarily the new card deployment.

The CEB concurs that whilst the system supports only low fare and purse values, it is a relatively unattractive target for criminals (although one would expect more public proof-of-concept demonstrations). In which case, the existing wide range of back-office detection and response mechanisms may balance the threat in the short term. However this balance may not be maintained when the national system is in place, as the higher ticket and purse values will make the system a far more attractive attack target. The remedial measures suggested by TNO could protect against some attack scenarios, however their practicality is not yet well proven and this should be investigated as a matter of urgency.

Migration readiness as defined by TNO is a very significant stage, as PTOs would need to have passed through a number of critical planning phases and committed significant investment in equipment and man-power to roll out the new infrastructure. Currently, the CEB can find no statistics or references that can reliably be used to verify or dispute whether the TNO two-year estimate for migration readiness is correct. To investigate this further would require detailed migration plans, which would also help to determine whether there should be a single trigger event for migration or rather progressive deployment as security problems increase in certain areas. There is a suggestion in the TNO report that technology needs to advance before fast and low-cost attacks become available to criminals and this may be a factor in the two-year estimate. However recent research would seem to contradict this view.

The CEB strongly recommends an earlier interim milestone referred to as the **Migration Planning Milestone,** set for January 2009 to coincide with the scheduled completion of the national roll-out for the current system. This is to ensure that from the start of nation-wide usage there is a state of preparedness for the migration to a higher level of card security. The migration plan should define all necessary activities, involved parties, budgets and technology. Providing open communication on progress towards the milestone may have a deterrent effect on attackers and the independent review of draft versions of the plan should provide added confidence that migration will succeed. It would require significant activity prior to the milestone, such as conducting a structured risk analysis, identifying the new card technology, defining infrastructure upgrades, selecting suppliers, arranging budgets and all other normal logistic, and project planning details. These steps, which are similar to those suggested by TNO, would stop short of any physical deployment. The final decision to start the migration could take place some time later, based on agreed processes and triggers (such as the measured level of fraud[4]) defined within the migration plan.

The CEB recognises that ideally a migration plan should be available earlier than the Migration Planning Milestone, but the programme of work would be very challenging and the transport companies should allow some time for consultation with members of the international expert community, as has been strongly recommended by a number of researchers. The last thing that anyone should want is a rushed and ill-informed decision on migration technology and then a poorly planned and executed deployment.

[4] Attributed to the card security issue

By the time the proposed Migration Planning Milestone is reached, there will also have been a reasonable period in which to collect the initial statistics relating to actual attack activity and to assess the effectiveness and practicality of introducing the remedial measures recommended in the TNO report. The earlier introduction of the remedial measures should be considered if the benefits justify the cost and effort, compared to waiting for the card migration solution.

The TNO report did not include proposals that could be guaranteed "future-proof" and so based on its own experience, the CEB would like to make additional suggestions for consideration during the migration planning activities. The card/algorithm migration should not be considered as a one-off activity, but an opportunity to prepare the system to cope with any future migration needs. The short-term recommendation is therefore to consider modularisation of the card and reader security solution so that multiple card/algorithms may be supported. This should not only help to future-proof the system, but would support the phased migration between card types. A benefit of phased migration is that existing cards may be allowed to reach the end of their normal lives before replacement, thereby reducing costs and inconvenience to customers.

For the long-term future, the modularisation approach might be extended further, as has happened in the mobile communications industry. The interface to the smart card authentication algorithm and its general functional requirements are defined and published in international standards, however the authentication algorithm itself is not standardised[5]. This permits companies to differentiate their security solutions whilst providing a framework that supports phased migration of algorithms. In the case of third generation (3G) mobile telephony, an international group of security experts published an open and freely available example algorithm, which is now widely adopted.

Whatever methods and experts are used to ensure that an appropriate replacement algorithm is selected for the OV-Chipkaart, it is equally important to implement the algorithm in a secure and attack resistant manner. Consideration should therefore be given to the independent evaluation[6] of the algorithm implementation within the smart card and system infrastructure.

Another general approach for future-proofing any system, is to pay particular attention to risk modelling. In fact the CEB would recommend in the strongest possible terms that rigorous application of risk assessment methodologies should be used when designing, deploying, operating and enhancing critical systems of this nature. In the case of the OV-Chipkaart system, the risk assessment should include all card types in current use, including the Mifare Ultralite.

To conclude, the CEB would like to put the current problems into perspective by highlighting that smart card based systems used in other countries and industry sectors are successful and popular with customers, offering fast, flexible and convenient service. Organisations involved with ticketing are often able to reduce ticket production and distribution costs and identify/control (perhaps for the first time) the non-technical frauds that are possible with paper or magnetic stripe tickets. In introducing a smart card based transport system to harness these benefits, the major challenge is the deployment of the system, the business infrastructure and establishment of associated operational processes. The customer experience is also vital as a transport system is intended for the convenient "fast-flow" of passengers. Whilst the smart card and associated reader modules are critical elements in the system solution, they represent a relatively small part of the infrastructure deployment challenge and they should therefore be practical to replace as part of a phased migration process.

[5] Note the data encryption algorithms are standardised to support equipment interworking and roaming

[6] Preferably in accordance with international standards