



16-CSRE-004m

Jury report Dutch Cyber Security Research paper Award 2016

Background

In the last 20 years, the Dutch information security community grew from a handful of brilliant mathematicians to a large community of cybersecurity researchers with representatives from many of the technical and the social sciences. The Netherlands Organisation for Scientific Research (NWO) and the European Commission have provided over a hundred million Euros of funding in long term Dutch cybersecurity research. This has led to dozens of new businesses, hundreds of highly skilled employees in all major corporations, government departments and universities, and thousands of scientific publications and patents.

Role IIP-VV and Jury

In 2015 the public-private Dutch ICT Innovation Platform on Security and Privacy (IIP-VV) decided to introduce a new and prestigious prize for the best recent Dutch scientific cybersecurity research paper. The ICT.OPEN conference is an excellent place with the right audience to announce the winner! The first DCSRA-winner was announced at the ICT.OPEN 2015 conference. The organization committee of IIP-VV board members designed the nomination and assessment process and came up with a long list of possible (foreign) jury members. Based on this list NWO Physical Sciences composed a jury with the task to assess eligible, i.e. recent non-commercial scientific cybersecurity research papers, received as a result of a call for nominations. For the second time the Jury, under technical chairmanship of NWO, consists of three well-respected scientists in the cybersecurity field. And like in the previous contest, again 16 papers were received, out of which the Jury selected the Top Five.

'Highlights of Dutch Cyber Security Research'

The cybersecurity track, part of the ICT.OPEN 2016 conference program, is scheduled March 22nd 2016 and is titled '*New challenges in cybersecurity and privacy*'. The objective of this track is to demonstrate the progress and achievements in the execution of recent cybersecurity research.

Within the session '*Highlights of Cyber Security Research*' the primary authors of the Top Five papers (as far as available) present their publications. Each presenter receives a "Dutch Cyber Security Research paper Award" certificate, signed by Jury members from the chair.

The Jury not only selected the Top Five, but also determined which paper ranks as the very best out of this set of five. The main author and presenter of the winning paper receives the "Dutch Cyber Security Research best paper Award" from DCSRA 2016 jurymember Prof. Dr. Srdjan Capkun (ETH Zurich, Switzerland) and a special bonus donated by IBM. This bonus, a € 500,- cheque, to honour the team of authors responsible for the best research paper, is personally handed over by the IBM Director Security Europe, Johan Arts.

**ICT.OPEN 2016 is a conference organized by NWO, STW and IPN.
The Dutch Cyber Security Research best paper Award 2016 is sponsored by IBM.**





16-CSRE-004m

Top Five Paper Reports

Practical Context-Sensitive CFI (paper 12)

Victor van der Veen, Dennis Andriessse, Enes Gökta, Ben Gras, Lionel Sambuc, Asia Slowinska, Herbert Bos, Cristiano Giuffrida

Main Author: Victor van der Veen, Presenter: Dennis Andriessse, published at CCS 2015

Motivation by Herbert Bos:

The Practical Context-Sensitive CFI paper by Victor van der Veen added a whole new research direction toward stopping highly advanced attacks. More precisely, the paper adds a novel chapter to Control Flow Integrity (CFI)—one of the few remaining lines of defence against really advanced attacks (and perhaps the most promising chapter). What is especially important for this award: Victor did so by refuting an old and crucial assumption that underpinned all CFI work since 2005. The paper was published at CCS, often considered to be the #1 security conference. The seminal work on CFI is now over a decade old. The idea is simple: CFI wants to guarantee that every branch in the program can jump only to code that corresponds to what the programmer intended. In other words, the target of the branch has to follow the original control flow graph of the program. This prevents an attacker from changing the control flow of the program in such a way that it executes instructions of the attacker's choosing—frustrating almost all exploitation attempts. CFI is regarded as one of the strongest and most promising approaches to thwart state-of-the-art exploits. Since 2005, the number of follow-up papers that try to improve CFI is almost beyond counting. Even today, every top venue delivers new solutions. Typically, they strive to make CFI either faster and more practical, or more fine-grained and thus more secure. Unfortunately, most CFI solutions presented at one conference get bypassed with new attacks at the next. The reason is that while these solutions change the approach to CFI somewhat, they follow the original general idea: at every branch they check whether the target address is in the original control flow graph when considering each branch in isolation. This leads to schemes that are overly-permissive. For instance, if a method `foo()` takes a callback as an argument, we allow `foo()` to call any function that could (potentially) be provided as a callback argument. Likewise, if a `printf()` function is called from 50 locations in the program, it means that we allow a return from that function to jump to any of these 50 call sites. That is crazy! In reality we should only allow `printf()` to return to its actual caller and not all potential callers. Similarly, if `foo()` was called from a context where only callback C1 can be the argument, we should not allow C2, C3, C4, etc. In other words, we should not look at a branch in isolation, but also take into account the context. A simple analogy is a luggage belt in an airport. When offloading an airplane, the luggage handler can, in principle, target every possible belt. However, this would lead to unhappy travellers. In reality, the plane from New York should target only belt #5, and the plane from Kuala Lumpur only belt #17, and so on. In other words, when performing the check, we should take into account the context—where did it come from? Even the original CFI paper admitted that context sensitivity would be ideal, but dismissed it as impractical: it would be way too expensive to track and check the entire context all the time. Also, it would be too hard to implement. All CFI papers since have similarly disregarded the context. What the paper by Victor and his colleagues shows is that on modern processors it is possible and practical to take the context into account when making CFI decisions. By making use of the CPU's abilities to store a short trace, lightweight monitoring of the context is now possible. To wit: the approach has as little as 3% overhead on the SPEC benchmark! There are several other innovations in the paper (e.g., JIT static analysis), but I believe that the fact that, for the first time, we can add the context to CFI policies is a watershed. It will enable a wave of future papers with better and more efficient context sensitive analysis that raise the bar for attackers. It is my belief that Practical Context-Sensitive CFI represents a major qualitative step forward in a very important field.



16-CSRE-004m

Jury's assessment:

Control Flow Integrity is a very important technique in improving software security by preventing attackers to change the control flow of a computer program. The submitted paper is an enhancement to state-of-the-art CFI and runtime attacks, which is a very relevant topic. It shows how limitations of current approaches can be overcome by taking the context into account. The researchers have chosen a solid methodology with extensive validation. This is done by focusing on checking software paths to sensitive program states. The only weakness is that applying this in practice may require some specific expertise. In short: a very good technical paper about research executed in collaboration with a company that is doing very promising work.

Post-Mortem of a Zombie: Conficker Cleanup After Six Years (paper 07)

Hadi Asghari, Michael Ciere, and Michel J.G. van Eeten

Main Author: Hadi Asghari, Presenter: Michael Ciere, published at the 24th USENIX Security Symposium

Motivation by Michel van Eeten:

Much work on botnet mitigation is on short term effects of takedown efforts. But mitigation is a long haul effort that has been understudied. This interdisciplinary paper undertakes the painstaking and innovative work to transform noisy sinkhole data into a statistically trustworthy set of parameters amenable to social science models. They allow us to better understand the fight against botnets over the long haul.

Jury's assessment:

Conficker was first detected in 2008, and is one of the most widespread botnets that infects machines running Windows OS. It has been six years since the botnet was sinkholed in a national cleanup of infected end user machines. The paper describes an excellent and detailed longitudinal statistical analysis of a large dataset obtained through the sinkhole of the Conficker worm, with some conclusions that are useful for a broader audience as well. This required an extensive amount of work. The governance, how ISP's can mitigate botnets, was studied. Unfortunately the central research question "Do national initiatives help in controlling a botnet" cannot be completely answered from the data, since apparently ISP's have not given eradicating Conficker a high priority. Overall a great paper to read with lots of data.

Ciphertext-only Cryptanalysis on Hardened Mifare Classic Cards (paper 09)

Carlo Meijer, Roel Verdult

Main Author and Presenter: Carlo Meijer, published at CCS 2015

Motivation by Lejla Batina

The breaking of the Mifare Classic chipcard in 2008 by Nijmegen's Digital Security group made headlines worldwide, and acted as a wake-up call for the security industry. Roel Verdult has been a leading researcher in this work. As a result, many organisations decided to abandon the Mifare chip altogether, and migrate to a chip with a higher protection level. But also, under pressure of organisations that could not migrate so easily, the Mifare implementation was tweaked and improved as much as possible within the standard. This worked for instance for Translink Systems, the organisation that operates the OV-chipcard in the Netherlands. Because of these tweaks the OV-chip fraud levels have been almost zero. Several researchers in Nijmegen have been involved in these improvements. Within this improvement team it was known that one particular vulnerability of the Mifare chip was inherent, and unfixable. This involved a problem that lies deep in the standard, but is difficult to exploit. During the improvement ('tweak') phase, it was treated as an accepted-risk. The current paper shows how to exploit this remaining, non-trivial vulnerability. It is the definitive death-blow to the Mifare chip. The paper is written



16-CSRE-004m

by Roel Verdult, together with Carlo Meijer, at the time a Kerckhoffs master student in Eindhoven (but now a PhD student in Nijmegen). This work is part of his master thesis.

This is a perfect show-case of how responsible disclosure can work very well.

Jury's assessment:

This paper was presented at a top conference and disproves the industry claim that the Mifare Classic contactless smartcard offers adequate security. Many attacks on Mifare Classic were already known; however they require access to the card reader or eavesdropping card - reader communication. The paper describes a card-only attack with consumer grade hardware that only requires a few minutes of wireless interaction with the card, in an environment without risk of camera detection. It is a very significant, valuable study. The technical execution is solid with a detailed analysis. Radboud University is known for their earlier discovery of weaknesses in the Mifare Classic card. There is no doubt that this paper had important impact on the Dutch industry (NXP). If this paper has the consequence that Mifare Classic is finally taken out of service after a responsible disclosure, this is a positive result. But Mifare Classic for several reasons should have been taken out of service already many years ago.

A measurement study of DNSSEC misconfigurations (paper 16)

Niels L. M. van Adrichem, Norbert Blenn, Antonio Reyes Lúa, Xin Wang, Muhammad Wasif, Ficky Fatturrahman and Fernando A. Kuipers

Main Author and Presenter: Niels L. M. van Adrichem, published in Security Informatics

Motivation by Jan van den Berg

This journal paper is an extended version of a conference publication that received a nomination for the best paper award, namely in the IEEE Joint Intelligence and Security Conference 2014 (IEEE ISI & EISIC). In the paper, the deployment of DNSSEC, an extension that aims to secure one of the most important technologies in the Internet, the DNS system, that offers protection against spoofing of DNS data by providing origin authentication, ensuring data integrity and authentication of non-existence by using public-key cryptography, is measured. The most innovative parts, are lying in the massive measurement and analysis of DNSSEC misconfigurations, which are categorized and explained for their possible causes. Additionally, the effects of misconfigurations on the reachability of a zone's network are evaluated where results show that out of the 4% of domains that show misconfigurations, almost 75% were unreachable from a DNSSEC-aware resolver. This illustrates that although the authorities of a domain may think their DNS is secured, it is in fact not. Worse still, misconfigured domains are at risk of being unreachable from the clients who care about and implement DNSSEC verification, while the publisher may remain unaware of the error and its consequences.

Jury's assessment:

Domain Name System Security Extensions (DNSSEC) offer protection against illegitimately falsifying data stored in DNS. The submitted paper describes an experimental study that analyzes failures of DNSSEC in practice, like demonstrating the effect of misconfigurations, for instance resulting in internet users unable to reach an online service. It is very insightful to study these issues. The paper also identifies some new failure modes. The methodology is solid and the results are helpful. The paper lacks a discussion of tools and/or methods to solve the problems identified.



16-CSRE-004m

Protocol State Fuzzing of TLS Implementations (paper 10)

Joeri de Ruiter, Erik Poll

Main Author: Joeri de Ruiter, Presenter: Erik Poll, published at the 24th USENIX Security Symposium

Motivation by Lejla Batina:

There has been a lot of research into the analysis of abstract security protocols (resulting in tools such as ProVerif), but very little research into the analysis of actual implementations of security protocols. This is a pity, because even if a security protocol is in principle sound, implementations of it can be - and regularly are! - insecure due to software flaws. This PhD research of Joeri de Ruiter is one few exceptions: he developed a systematic approach to detect flaws in the program logic of implementations of TLS. The technique is automated and only relies on black box testing. He applied it to TLS though the technique can be applied to any protocol. He proved the value of the technique by using it to reveal flaws in three leading TLS implementations (OpenSSL, GnuTLS, and the Java Secure Socket Extension), all of which have been fixed by now.

Jury's assessment:

Transport Layer Security (TLS) protocol implementations are very important in today's internet security. The paper, presented at a top conference, uses an existing tool for black box analysis (state machine learning) techniques to recover the protocol state machine of commonly used implementation of TLS. Several new flaws were revealed, and it is also shown that several implementations have state machines which are more complex than needed. A clear conclusion is that state machines should be included in official protocol specifications to reduce implementation freedom. This approach has also been used for another security protocol (EMV). All in all, the Jury likes this paper very much for the solid scientific approach, the impact, the relevance, the excellent quality of the write-up and the conclusions with clear recommendations.

The Winner

Jury members, who individually ranked and collectively decided on the quality of the papers received, appreciated the response by the Dutch research community on the call for nominations, resulting in 16 paper nominations. Out of the Top Five as presented today, one paper deserves the predicate best Dutch cybersecurity research paper.

*Jury member Srdjan Capkun presented a special certificate to the main author of this paper Joeri de Ruiter, because the Jury chose his paper "**Protocol State Fuzzing of TLS Implementations**", co-authored by Erik Poll, as the best Dutch cybersecurity research paper 2016.*

Johan Arts, IBM Director Security Europe, offered a bonus cheque to the winner who is represented by co-author Erik Poll.

Jury members

- Prof. Dr. Srdjan Capkun (ETH Zurich, Switzerland)
- Dr. Wee Keong Ng (Nayang Techn. University, Singapore)
- Prof. Bart Preneel (CU Leuven, Belgium)
- Drs. Jan Piet Barthel (NWO, IIP-VV/CSRE-Platform, The Hague, The Netherlands) – chair