

Research Statement

Tobias Fiebig

Abstract

For my research I methodologically focus on obtaining a *practical* understanding of relevant *real-world* challenges using *empirical methods*. In the past I used this approach to obtain a deeper understanding of the problem space of security misconfigurations in Internet services. In this context, misconfigured services are those that are exploitable due to human error during their deployment.

Following up on this research I currently pursue various directions of further work. Current and future research includes: (i) Obtaining a deeper understanding of the human factor in security misconfigurations, (ii) Further research on IPv6 recognisance and privacy considerations in DNS, where I already lead an international collaboration on the matter with researchers from the University of California, Santa Barbara, (iii) Projects on the security considerations during protocol design, and the impact of network protocol design on implementation security, and, (iv) Research on distributed IoT environments like distributed sensor networks, as signified by my under-submission grant proposal for a secure update delivery infrastructure for autonomous cars. Furthermore, I plan to transfer my earlier research results on wide-spread service misconfiguration in Internet services to the domain of distributed sensor networks. Due to their versatile requirements and the rise of IoT and the Industry 4.0 they pose a set of important research areas. Research areas with open challenges include the basic system security of smart sensors, especially in the context of *simple* yet common vulnerabilities, network security in sensor networks, and, Big Data systems using learning algorithms to evaluate data collected from sensor networks. In addition, work on these systems can also be actively integrated into teaching duties in accordance with the goals presented in my teaching statement.

In this document, I provide a summary of my dissertation and an overview of my previous grant success in Section 1. Subsequently, I discuss my current and future research in Section 2. The document concludes with a list of my peer-reviewed and currently under-submission scientific articles and their abstracts in Section 4.

Contents

1	Prior Research	2
1.1	Dissertation	2
1.2	ENZEVALOS: Industry Integration, Standardization & Public Use	3
2	Current and Future Research	3
2.1	Human Factors in Operations	3
2.2	IPv6 Recognisance and DNS Security and Privacy	4
2.3	Security in Protocol Design	4
2.4	AutoSEKURE: Secure Infrastrucutre for Autonomous Cars	4
2.5	Distributed Sensor Networks and Adaptive Big Data Systems	5
3	Summary	6
4	Scientific Articles	6
4.1	Published	6
4.2	Under Submission	8

1 Prior Research

So far, the core-focus of my research has been on the security of distributed systems, identifying systematic issues in their deployment and design. Specifically, I performed an interdisciplinary empirical investigation of human error in the design, implementation, and operation of distributed systems. I find, that the most crucial issue for the security of modern distributed systems do not stem from *complex* attacks, but instead in a neglect of *due dilligence* by the engineers implementing and operating these systems.

My empirical research on attacks and their prevalence left me with a set of interesting *practical challenges*. Based on these, I designed an industry focused third-party funds project, for which I successfully obtained funding from the BMBF (Bundesministerium Bildung und Forschung). This section first provides a summary of my dissertation and a summary of the key-insights for the practical application of my work. Subsequently, I detail how the findings that are summarized in my thesis lead to applied work in the context of third-party funds.

1.1 Dissertation

My thesis is motivated by the dispersion between the focus of IT security research on complex attacks and a perceived dominance of *simple human errors* when it comes to the root cause of security incidents affecting billions of end-users. Indeed, the literature is full of sophisticated attacks to obtain confidential information from computer systems, compromise them, or prevent them from being used at all. Simultaneously, mitigations to these attacks are as well studied. Technically, many current attacks could be mitigated by deploying these techniques. In fact, there is a constant stream of new, complex techniques to ensure the confidentiality, integrity, and availability of data and systems. However, even the recent past has not been short of security incidents affecting billions of people. Yet, these incidents are usually neither enabled nor mitigated by these complex techniques. On the contrary, I find that these breaches are usually caused by something far more simple: Human error in deploying and running network services, e.g., delayed software updates, or, no authentication and authorization being configured even though it would have been available. We refer to these oversights in due dilligence as security misconfigurations.

Hence, I conducted an interdisciplinary empirical investigation of the nature of security misconfigurations. Specifically, I focus on: (i) What misconfiguration based attacks are, (ii) Which practical impact security misconfigurations have, (iii) How security misconfigurations can be measured and mitigated in today's as well as tomorrow's Internet, and (iv) Which possible root-causes for security misconfigurations I can find in general. Furthermore, I investigate if security misconfigurations are limited to network services, or if similar root causes lead to security vulnerabilities in data plane components in the context of SDN.

Initially looking at complex attacks, I find that they are commonly prevented by *easy to implement* technical mitigations. In contrast, misconfiguration based issues require a more demanding approach that is focused on the personnel operating Internet services. Patching humans is incredibly hard. Furthermore, an additional literature study indicates that major misconfiguration traps are introduced during the design of systems and protocols. A good design can prevent misconfigurations, while a bad design can lead to multiple, easily misconfigured implementations.

Current techniques for mitigating misconfigurations are focused on identifying and contacting affected operators, so they can take appropriate action and remove the misconfiguration. However, this process heavily relies on security scans of the whole Internet. While this is feasible with IPv4, the current Internet Protocol Version, a similar bruteforce approach is unfeasible for the larger address space of the upcoming IPv6. Hence, I develop and evaluate a new methodology that enables researchers to perform security scans of IPv6 connected hosts. In addition, I conduct an analysis of security paradigms in new data plane components introduced to accommodate the developing cloud landscapes in modern data centers. Similar to my analysis of protocol design, I find that already the concepts themselves introduce vast attack surfaces.

In summary I find that the major objective for the future are: (i) The introduction and enforcement of *due diligence* by engineers in the design, implementation, and operation of Internet services and distributed systems, (ii) The identification and creation of systems that are—by default—resilient against human error, (iii) a further investigation of the human factor in the operation of distributed systems, and (iv) the extension of my Investigation to systems beyond Internet services and SDN systems, i.e., IoT (Internet of Things), Industry 4.0, and Cognitive Systems.

1.2 ENZEVALOS: Industry Integration, Standardization & Public Use

Motivated by the need to create protocols, services, and techniques that are resilient to misconfiguration and other security issues like widespread surveillance, I designed and obtained funding for project ENZEVALOS together with Prof. Dr. Roth from FU Berlin. The objective of ENZEVALOS is finding methods that increase the widespread use of end-to-end encryption for end-users. This reduces the impact of server side misconfiguration, as sensitive user data is encrypted with keys only accessible by the end-users. With his research interests being mostly on usability in security [Roth et al., 2005], he planned to focus on usability aspects of the public key-distribution problem [Roth et al., 2005]. To approach this, we envisioned a distributed provider-signature assisted trust model.

In this context I contributed the solid foundation in standards, which are necessary to have a system accepted and adopted by the industry. In planning the research tasks for ENZEVALOS, I consequently realized this requirement. By integrating OpenXchange, Germanies leading supplier of enterprise level webmail systems, as associated partner, I established a direct interface to the industry. This enabled an efficient technology transfer of the research results into market ready products. In addition, I also planned to participate in the Internet Engineering Task Force (IETF) within the context of ENZEVALOS. The IETF is, even before the International Organization for Standardization (ISO) and the International Telecommunication Union (ITU) the relevant standards body for the Internet. IETF documents, called Requests For Comments (RFCs) have become the basis for interoperable protocols on the Internet. Accordingly, ENZEVALOS aims at contributing an RFC for the designed provider trust system. Attending the meetings of the IETF offered me direct contact to representatives of all major industry players, and allowed me to build a network of contacts with whom I could integrate industry requirements and research results.

2 Current and Future Research

2.1 Human Factors in Operations

My previous work signifies that human errors in operations are a crucial problem. However, my work also indicated further research directions in the context of security misconfiguration. In my further research, I will first focus on the introduction and enforcement of *due diligence* by engineers. To approach this issue, we have to first conduct further research on the human factor in the operation of distributed systems. This aspect of IT security has been mostly neglected in the past. Security and usability research usually focused on the end-users, e.g., [Redmiles et al., 2016b, Redmiles et al., 2016a]. While more recent work [Acar et al., 2016] starts to investigate programmers as a human factor in security, operation personell is still insufficiently considered in the research community.

To accomplish this, I currently steer a project aimed at understanding the perspective of operators on misconfiguration. We currently apply a multi-step human research approach (informal pre-interviews, focusgroups, structured interviews, questionair design, quantitative study) targeted at systemadministrators and operations personell. For this purpose I leverage connections to various operations related conference venues, in order to obtain access to a sufficient sample of operators. So far the project looks promising, and the participating student will already present her results at one of the major european operations venues. Furthermore, I am in contact with researchers from the Carnegie Mellon University who would like to collaborate on the dataset which we are currently obtaining.

2.2 IPv6 Recognisance and DNS Security and Privacy

The biggest challenge for the Internet as it exists today is the exhaustion of the legacy IPv4 pool [Richter et al., 2015]. IPv6 has been introduced as a successor to IPv4, but still lack behind in adoption. Nevertheless, IPv6 also introduces severe security challenges [Czyz et al., 2016, Chittimaneni et al., 2017]. In the context of security, exhaustive scans of the IPv4 address space have become an important tool for empirical research. Especially the large-scale detection of vulnerable IoT devices is severely hampered by IPv6, which can not be scanned exhaustively.

The work of me and my collaborators is the first step in solving this problem. I currently pursue these challenges in close collaboration with researchers from the University of California, Santa Barbara (UCSB). This also includes a three month research visit to UCSB. There I designed and deployed a measurement setup that allows us to conduct a more time-scaled investigation of IPv6 security. While we are still in early stages of this work, we were already approached by various major industry and NGO players that were interested in our research. We expect the basic methodological research to be completed in May 2017. Thereafter we plan to perform applied investigations of our method, focusing on the security of Machine-to-Machine communication/IoT on the IPv6 Internet.

2.3 Security in Protocol Design

Standardization and protocol definitions are a major challenge for designing secure systems. While it is important to establish standards that facilitate secure systems, the design of standards may actually lead to vulnerable systems on the Internet [Fiebig et al., 2016]. As a follow-up, I plan to investigate if and how the complexity of protocol grammars leads to implementations that are prone to memory access violations.

With the arising interconnected world, we will face a constant stream of new protocols being developed. Especially in the context of the IoT, Industry 4.0 and mobile sensor networks, we will see a countless number of new protocols. My current work demonstrates the foundations of this issue. Together with my collaborators from the Security in Telecommunications researchgroup (FG SECT) at TU Berlin I identify crucial vulnerabilities in virtual Switches. We demonstrate how protocol parsers in virtual switches become vulnerable due to the parsed protocol's design. In conjunction with the design of virtual switches this leads to severe weaknesses [Thimmaraju et al., 2016].

In a second work, which is currently under submission, we demonstrate that the underlying protocol grammar for a network application can be used to infer input rules using static analysis that improve the efficiency of state-of-the-art fuzzers. Hence it is imperative to investigate how protocols can be designed to not facilitate, but prevent security issues in their implementations.

2.4 AutoSEKURE: Secure Infrastrucutre for Autonomous Cars

Due to the constant support of my advisors, who allowed me to write grant proposals for my own ideas and concepts under their patronage I have already aquired one grant (ENZEVALOS, discussed earlier). However, Prof. Anja Feldmann, Ph.D. and Prof. Dr. Seifert also allowed me to build a second consortium for a project on secure network infrastructures for autonomous cars. Contrary to ENZEVALOS, AutoSEKURE is still under submission. We expect feedback from the fundingbody within the next couple of months.

The goal of this project is the design of a resilient infrastructure for the secure distribution of updates. However, in terms of autonomous automotive systems, user-acceptance is of ourmost importance to the success of any system. Hence, the consortium I brought together consists of TU Berlin as the technical academic partner and the Institute of Cognitive Science from the University of Osnabrück as the academic partner for user acceptance. I selected the partner from the University of Osnabrück due to their prior experiments in virtual reality based experiments on beheavioural acceptance of autonomous systems actions [Skulmowski et al., 2014].

To give this project the necessary connection to the Industry, I was able to use my personal network to aquire the MinteTronics GmbH as an additional partner for the project. The MineTronics GmbH is

a “hidden tiger”, providing engineering solutions for the mining and bulk resource generation industry. They are the first company to demonstrate an autonomous resource collection vehicle, but also provide solutions for measurement and communication systems in industrial environments. Hence, their role within AutoSECURE is the accompanying development of the hardware required to realize the research results from TU Berlin and University of Osnabrück. In addition, I was also able to integrate Deutsche Telekom AG as an associated partner. They plan to contribute their experience in deploying country-scale network infrastructure to the project.

2.5 Distributed Sensor Networks and Adaptive Big Data Systems

The early results of my IPv6 recognisance research already indicate that the security of IPv6 connected systems is a significant future challenge. An upcoming area that will be most affected by this is the security of distributed sensor networks. This also includes the security of the Big Data systems that are fed by these sensor networks. I see three major, interconnected challenges in this area: (a) The system security of single sensors, (b) The security of sensor networks, and, (c) Security and privacy in connected (big data) platforms, which may be compromised by crafted sensor data.

Sensor System Security: Currently, sensors (smart meters, smoke detectors, sun-sensors for solar panels etc.) for sensor networks are commonly realized on the Linux platform [Bassi et al., 2013], using semi-general-purpose hardware (ARM, MIPS). This makes these systems susceptible to a variety of attacks. However, even more so, the security of these systems is commonly tainted by misconfigurations or misconfiguration related vulnerabilities, e.g., an insufficiently protected telnet interface or weak default credentials [Pa et al., 2015]. However, the first step in securing sensor networks is securing the end-devices.

For this, I plan a two-sided approach:

1. Finding yet another vulnerability in yet another sensor is in itself not science. However, constantly testing new sensors for vulnerabilities is an important tool in understanding the development of embedded device security. Furthermore, performing a security test on a sensor is a viable tool to teach a student security engineering practices on a real-world subject (see my teachingstatement for further information).
2. With the problem space being similar to security misconfiguration in Internet services, I plan to adapt learnings from my ongoing research on misconfiguration. Techniques and best practices to reduce misconfiguration in Internet services are certainly adoptable to vulnerable embedded sensor products. In fact, a portion of devices considered misconfigured Internet services are, in fact, embedded devices. Furthermore, the root-cause analysis process already initiated for misconfiguration can also be applied to sensor systems vulnerable to simple exploits.

Sensor Network Security: The second important aspect in the context of distributed sensor networks is the actual network security on the sensor network. Here, confidentiality, integrity and availability have to be protected. This poses some unique challenges:

Confidentiality Confidentiality on the wire is commonly ensured using encryption. However, sensors are commonly produce small data samples (less than the block size of established cryptographic algorithms), which are frequently sent, located in a small range of possible values. This is one of the hardest corner cases for encryption algorithms, catering towards replay and chosen/known plain/cipher text attacks. Standard techniques to handle this, like padding, are usually not feasible due to limited entropy on embedded devices [Heninger et al., 2012].

Integrity Integrity in the context of sensor networks pertains not only to the transmitted sensor readings, but also to the integrity of the sensor’s operating system. While integrity suffers from the same cryptographic challenges we discussed for confidentiality, the latter can be easily compromised due to insufficiently protected management access [Pa et al., 2015], as already discussed in the context of device security. However, the counterpart to (single) device management is a sensor network’s backend, which possibly holds access opportunities to all sensors in a network.

Availability Depending on the specific use case, availability in sensor networks may hold different interpretations. A sensor that is only periodically polled may not require uninterrupted network access. In addition, a sensor, like a smart meter, for data that is not required in (near) realtime may use caching to buffer small network outages. However, a sensor which feeds into a live monitoring and alerting system, for example the pressure sensors for a power plant’s steam tank, has strict requirements for availability. In traditional scenarios, this would have been accomplished by analog safety measures. With the rise of “Industry 4.0”, however, vendors are turning towards Ethernet and TCP/IP as transport layers.

These challenges are, especially for above-metro-level sensor networks, similar to those which I already plan to address in the AutoSEKURE grant proposal. However, as seen above, the context of sensor networks also adds interesting new research opportunities, which I will address in the future.

Big Data platforms for sensor networks: Another aspect of sensor networks are analysis systems. These are already in active use for the large-scale analysis of (non) public measurement data. The most prominent example for this are currently pre-crime efforts which are heavily based on an analysis of prior crime patterns, but can also be augmented with other sensory data. This however enables attackers to potentially feed crafted information to the system. So far, there has been no in-depth analysis on the attack opportunities induced by malicious/crafted inputs to such systems.

However, attacks on learning based data analysis systems may have severe impact, depending on the systems that are being monitored. For example, a system monitoring the power consumption in a smart grid, trained to make decisions on maintaining grid stability may be attacked. By inducing data that indicates oversaturation, the system reduces energy production, just while a spike occurs. With smart-meters becoming wide-spread, access to the sensors used to feed such systems becomes readily available for attackers. This also underlines that an investigation of the security of these systems is imperative.

Similarly, also low-cost attacks may be feasible. Attackers may try to deceive sensor networks to assume a certain state by applying crafted information to a selected set of sensors. Such attacks are certainly possible, and should be the matter of further research.

3 Summary

In summary, my research background expands between low level system security in virtualized network architectures, mobile security, and an empirical risk assessment of misconfiguration as presented in my thesis. I have identified and pursue several further research directions that follow from my earlier works. In addition, I have already demonstrated how the *practical* implications of my work can be used to obtain research grants that enable me to convert my ideas to industry adopted solutions. Following this concept, I also initiated a grant proposal for a resilient data distribution platform for autonomous cars, which is currently under submission. The design proposed in this draft could also be used for the distribution and recovery of sensor information in the context of the Internet of Things. In addition to expanding my current work on security misconfiguration, I want to transfer results from this domain to that of distributed sensor networks. I identified several areas where methodology and results from my prior work can be applied to improve security. Furthermore, to fund my future research into this topic, I am currently in the first stages of drafting a grant proposal on this matter.

4 Scientific Articles

4.1 Published

“**Something From Nothing (There): Collecting Global IPv6 Datasets From DNS**”, Tobias Fiebig, Kevin Borgolte, Shuang Hao, Christopher Kruegel, Giovanni Vigna, Passive and Active Measurement: 18th International Conference, 2017.

Abstract— Current large-scale IPv6 studies mostly rely on non-public datasets, as most public datasets are domain specific. For instance, traceroute-based datasets are biased toward network equipment. In

this paper, we present a new methodology to collect IPv6 address datasets that does not require access to restricted network vantage points. We collect a new dataset spanning more than 5.8 million IPv6 addresses by exploiting DNS' denial of existence semantics (NXDOMAIN). This paper documents our efforts in obtaining new datasets of allocated IPv6 addresses, so others can avoid the obstacles we encountered.

“A One-Year Perspective on Exposed In-memory Key-Value Stores”, Tobias Fiebig, Anja Feldmann, Matthias Petschick, Proceedings of the 2016 ACM Workshop on Automated Decision Making for Active Cyber Defense, 2016.

Abstract— Today's highly-scalable low-latency Web services rely on in-memory key-value stores. While they are essential to improve Web service performance they should not be exposed to the Internet. Security problems range from data leakage to remote code execution. In this paper we use a year long data set of exposed Redis and memcached instances to highlight the magnitude (about 200K) of the problem, document new transitive attacks, and explore misconfiguration patterns. We find that the number of exposed instances is constantly on the rise and that even severe problems only lead to temporal decreases. However, by correlating misconfiguration patterns we can explain significant changes in the number of exposed systems.

“SoK: An Analysis of Protocol Design: Avoiding Traps for Implementation and Deployment”, Tobias Fiebig, Franziska Lichtblau, Florian Streibelt, Thorben Krüger, Pieter Lexis, Randy Bush, Anja Feldmann, Technical Report arXiv:1610.05531

Abstract— Today's Internet utilizes a multitude of different protocols. While some of these protocols were first implemented and used and later documented, other were first specified and then implemented. Regardless of how protocols came to be, their definitions can contain traps that lead to insecure implementations or deployments. A classical example is insufficiently strict authentication requirements in a protocol specification. The resulting misconfigurations, i.e., not enabling strong authentication, are common root causes for Internet security incidents. Indeed, Internet protocols have been commonly designed without security in mind which leads to a multitude of misconfiguration traps. While this is slowly changing, to strict security considerations can have a similarly bad effect. Due to complex implementations and insufficient documentation, security features may remain unused, leaving deployments vulnerable.

In this paper we provide a systematization of the security traps found in common Internet protocols. By separating protocols in four classes we identify major factors that lead to common security traps. These insights together with observations about enduser centric usability and security by default are then used to derive recommendations for improving existing and designing new protocols—without such security sensitive traps for operators, implementors and users.

“Analyzing End-Users' Knowledge and Feelings Surrounding Smartphone Security and Privacy”, Lydia Kraus, Tobias Fiebig, Viktor Miruchna, Sebastian Möller, Asaf Shabtai, Security & Privacy Workshops 2015 - Mobile Security Technologies (MoST), 2015.

Abstract— Along with the significant growth in the popularity of smartphones and the number of available mobile applications, the amount of threats that harm users or compromise their privacy has dramatically increased. The mobile security research community constantly uncovers new threats and develops associated mitigations. Recently, there is an increasing interest in the human factors and various studies investigated user-aspects in the implementation of security mechanisms as well as users' perception of threats. In this paper we present a qualitative study on end-users' knowledge and perceptions of threats and mitigations on mobile devices. Moreover, we identify feelings surrounding smartphone security and privacy. We interpret these feelings in the context of basic psychological need fulfillment. Our findings suggest that so-far little considered aspects of why end-users do not utilize mitigations reside in the need fulfillment plane, and not only in the conflict of usability and security. Following these findings we give examples of how developers of mitigations could ensure that these mitigations are actually adopted by end-users.

“Security Impact of High Resolution Smartphone Cameras”, Tobias Fiebig, Jan Krissler and Ronny Hänsch, 8th USENIX Workshop on Offensive Technologies (WOOT 14), 2014.

Abstract— Nearly every modern mobile device includes two cameras. With advances in technology the resolution of these sensors has constantly increased. While this development provides great convenience for users, for example with video-telephony or as dedicated camera replacement, the security implications of including high resolution cameras on such devices has yet to be considered in greater detail. With this paper we demonstrate that an attacker may abuse the cameras in modern smartphones to extract valuable information from a victim. First, we consider exploiting a front-facing camera to capture a user’s keystrokes. By observing facial reflections, it is possible to capture user input with the camera. Subsequently, individual keystrokes can be extracted from the images acquired with the camera. Furthermore, we demonstrate that these cameras can be used by an attacker to extract and forge the fingerprints of a victim. This enables an attacker to perform a wide range of malicious actions, including authentication bypass on modern biometric systems and falsely implicating a person by planting fingerprints in a crime scene. Finally, we introduce several mitigation strategies for the identified threats.

“A metric for the evaluation and comparison of keylogger performance”, Tobias Fiebig, Janis Daniševskis, Marta Piekarska, 7th USENIX Workshop on Cyber Security Experimentation and Test (CSET 14), 2014.

Abstract— In the field of IT security the development of Proof of Concept (PoC) implementations is a commonly accepted method to determine the exploitability of an identified weakness. Most security issues provide a rather straightforward method of asserting the PoCs efficiency. That is, it either works or it does not. Hence, data gathering and exfiltration techniques usually remain in a position where the viability has to be empirically verified. One of these cases are mobile device keyloggers, which only recently have been starting to exploit side-channels to infer heuristic information on a user’s input. With this introduction of side channels exploiting heuristic information the performance of a keylogger may no longer be described with “it works and gathered what was typed”. Instead, the viability of the keylogger has to be assessed based on various typing speeds, user input styles and many metrics more as documented in this paper. The authors of this document provide a survey of the required metrics and features. Furthermore, they have developed a framework to assess the performance of a keylogger. This paper provides the documentation on how such a study can be conducted, while the required source code is shared online.

4.2 Under Submission

“Enumerating Active IPv6 Addresses via DNSSEC-signed Reverse Zones”, Kevin Borgolte, Shuang Hao, Tobias Fiebig, Christopher Kruegel, Giovanni Vigna

Abstract— Security research has made extensive use of exhaustive Internet-wide scans over the recent years, as they can provide significant insights into the overall state of security of the Internet, and ZMap made scanning the entire IPv4 address space practical. However, the IPv4 address space is exhausted, and a switch to IPv6, the only accepted long-term solution, is inevitable. In turn, to better understand the security of devices connected to the Internet, including in particular Internet of Things devices, it is imperative to include IPv6 addresses in security evaluations and scans. Unfortunately, it is practically infeasible to iterate through the entire IPv6 address space, as it is 296 times larger than the IPv4 address space. Therefore, enumeration of active hosts prior to scanning is necessary. Without it, we will be unable to investigate the overall security of Internet-connected devices in the future. Recently, Fiebig et al. suggested a new method to enumerate IPv6 hosts via reverse DNS pointers. We find that their technique can be easily mitigated. In fact, by now, more than 60% of the hosts found by Fiebig et al. can no longer be enumerated due to deployed mitigations.

In this paper, we introduce a novel technique to enumerate an active part of the IPv6 address space by walking DNSSEC-signed IPv6 reverse zones. Subsequently, by scanning the enumerated addresses, we

uncover significant security problems: the exposure of sensitive data, and incorrectly controlled access to hosts, such as access to routing infrastructure via administrative interfaces, all of which were accessible via IPv6. Furthermore, from our analysis of the differences between accessing dual-stack hosts via IPv6 and IPv4, we hypothesize that the root cause is that machines automatically and by default take on globally routable IPv6 addresses.

This is a practice the affected system administrators appear unaware of, as the respective services are almost always properly protected from unauthorized access via IPv4. Our findings indicate (i) that enumerating active IPv6 hosts is practical without a preferential network position contrary to common belief, (ii) that the security of active IPv6 hosts is currently still lagging behind the security state of IPv4 hosts, and (iii) that unintended IPv6 connectivity is a major security issue for unaware system administrators.

“Virtual Switches: Compromising Cloud Systems via the Data Plane”, Kashyap Thimmaraju, Bhargava Shastry, Tobias Fiebig, Felicitas Hetzelt, Jean-Pierre Seifert, Anja Feldmann, Stefan Schmid

Abstract— Virtual switches are a crucial component of cloud operating systems to interconnect virtual machines in a flexible manner. However, in this paper we demonstrate that they also introduce new attack surfaces. We find that these are caused by: (i) The co-location of data-plane components with the hypervisor’s kernel, (ii) The logical centralization of such networks (e.g., OpenStack or SDN), (iii) The presence of virtual switches on the network’s edge where they are directly exposed to attackers’ crafted input, and (iv) Extended protocol parsing.

We showcase the feasibility of such attacks on Open vSwitch, a popular open-source virtual switch implementation. By analyzing code-paths for packet parsing, we quickly uncover vulnerabilities using off-the-shelf analysis tools. These vulnerabilities let us compromise an OpenStack cloud within minutes. Furthermore, we analyze the applicability of our findings on technologies beyond virtual switches and find similar issues there.

We conclude the paper with an analysis on the applicability of various mitigation techniques, like ASLR, PIEs, and unconditional stack canaries as well as architectural changes. We find that these techniques limit virtual switches’ attack surface, but, that they are not well adopted. In summary, our work demonstrates that we have to reconsider existing threat models for virtualized data plane technologies.

“Static Program Analysis as a Fuzzing Aid”, Bhargava Shastry, Markus Leutner, Tobias Fiebig, Kashyap Thimmaraju, Fabian Yamaguchi, Konrad Rieck, Stefan Schmid, Jean-Pierre Seifert, Anja Feldmann

Abstract— Fuzz testing is an effective and scalable technique to perform software security assessments. Yet, contemporary fuzzers fall short of thoroughly testing applications with a high degree of control-flow diversity, such as firewalls, and network packet analyzers. In this paper, we demonstrate how static program analysis can be used to more effectively guide fuzzing by augmenting existing program models maintained by the fuzzer. Based on the insight that code patterns reflect the data format of inputs processed by a program, we automatically construct an *input dictionary* by statically analyzing program control and data flow. Our analysis is performed before fuzzing commences, and the input dictionary is supplied to an off-the-shelf fuzzer to influence input generation. Evaluations show that our technique not only increases test coverage by 10–15% over baseline fuzzers such as *afl*, but also reduces the time required to expose vulnerabilities by up to an order of magnitude. As a case study, we have evaluated our approach on two classes of network applications: nDPI, a deep packet inspection library, and tcpdump, a network packet analyzer. Using our approach, we have uncovered 15 zero-day vulnerabilities in the evaluated software that were not found by stand-alone fuzzers. Our work not only provides a practical means to more effectively conduct security evaluations but also demonstrates that the synergy between program analysis and testing can be exploited for a better outcome.

References

- [Acar et al., 2016] Acar, Y., Backes, M., Fahl, S., Kim, D., Mazurek, M. L., and Stransky, C. (2016). You get where you’re looking for: The impact of information sources on code security. In *Proc. IEEE Security & Privacy (S&P)*, pages 289–305.
- [Bassi et al., 2013] Bassi, A., Bauer, M., Fiedler, M., Kramp, T., Van Kranenburg, R., Lange, S., and Meissner, S. (2013). Enabling things to talk. *Designing IoT solutions with the IoT architectural reference model*, pages 163–211.
- [Chittimaneni et al., 2017] Chittimaneni, K., Kaeo, M., and Vyncke, E. (2017). draft-ietf-opsec-v6-10: Operational security considerations for ipv6 networks.
- [Czyz et al., 2016] Czyz, J., Luckie, M., Allman, M., and Bailey, M. (2016). Don’t forget to lock the back door! a characterization of ipv6 network security policy. In *Proc. Internet Society Symposium on Network and Distributed System Security (NDSS)*.
- [Fiebig et al., 2016] Fiebig, T., Lichtblau, F., Streibelt, F., Krueger, T., Lexis, P., Bush, R., and Feldmann, A. (2016). Sok: An analysis of protocol design: Avoiding traps for implementation and deployment. *arXiv preprint arXiv:1610.05531*.
- [Heninger et al., 2012] Heninger, N., Durumeric, Z., Wustrow, E., and Halderman, J. A. (2012). Mining your ps and qs: Detection of widespread weak keys in network devices. In *Proc. Usenix Security Symp.*, volume 8.
- [Pa et al., 2015] Pa, Y. M. P., Suzuki, S., Yoshioka, K., Matsumoto, T., Kasama, T., and Rossow, C. (2015). Iotpot: Analysing the rise of iot compromises. In *Proc. USENIX Workshop on Offensive Technologies (WOOT)*.
- [Redmiles et al., 2016a] Redmiles, E. M., Kross, S., and Mazurek, M. L. (2016a). How i learned to be secure: a census-representative survey of security advice sources and behavior. In *Proc. ACM Conference on Computer and Communications Security (CCS)*, pages 666–677.
- [Redmiles et al., 2016b] Redmiles, E. M., Malone, A., and Mazurek, M. L. (2016b). I think they’re trying to tell me something: Advice sources and selection for digital security. *Proc. IEEE Security & Privacy (S&P)*.
- [Richter et al., 2015] Richter, P., Allman, M., Bush, R., and Paxson, V. (2015). A primer on ipv4 scarcity. *ACM Computer Communication Review (CCR)*, 45(2):21–31.
- [Roth et al., 2005] Roth, V., Straub, T., and Richter, K. (2005). Security and usability engineering with particular attention to electronic mail. *International Journal of Human-Computer Studies*, 63(1):51–73.
- [Skulmowski et al., 2014] Skulmowski, A., Bunge, A., Kaspar, K., and Pipa, G. (2014). Forced-choice decision-making in modified trolley dilemma situations: a virtual reality and eye tracking study. *Frontiers in behavioral neuroscience*, 8:426.
- [Thimmaraju et al., 2016] Thimmaraju, K., Shastry, B., Fiebig, T., Hetzelt, F., Seifert, J.-P., Feldmann, A., and Schmid, S. (2016). Reigns to the cloud: Compromising cloud systems via the data plane. *arXiv preprint arXiv:1610.08717*.