SEVENTH FRAMEWORK PROGRAMME

Information & Communication Technologies
Trustworthy ICT

NETWORK OF EXCELLENCE

A European Network of Excellence in Managing Threats and Vulnerabilities in the Future Internet: *Europe for the World* [†]

Deliverable D4.2: Second Report on Threats on the Future Internet and Research Roadmap

Abstract: This deliverable presents an overview of current and emerging threats identified by the three working groups at the end of the second year of the project. In addition, this deliverable contains the updated version of the research roadmap in the area of System Security.

Contractual Date of Delivery	August 2012
Actual Date of Delivery	September 2012
Deliverable Dissemination Level	Public
Editor	Davide Balzarotti
Contributors	All *SysSec* partners
Quality Assurance	Magnus Almgren, Stefano Zanero

[†]The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 257007.

The *SysSec* consortium consists of:

FORTH-ICS	Coordinator	Greece
Politecnico Di Milano	Principal Contractor	Italy
Vrije Universiteit Amsterdam	Principal Contractor	The Netherlands
Institut Eurécom	Principal Contractor	France
IICT-BAS	Principal Contractor	Bulgaria
Technical University of Vienna	Principal Contractor	Austria
Chalmers University	Principal Contractor	Sweden
TUBITAK-BILGEM	Principal Contractor	Turkey

Contents

Acknowledgments

A number of researchers and external experts contributed to the discussion and refinement of the list of threats and to the research roadmap presented in this document. We would like to thank all the members of the Industrial Advisory Board, the 56 external members of the three working group mailing lists, and all the participants of the June face-to-face working group meeting. Their positive discussion and invaluable feedback was very important to help us define the content of this deliverable.

In particular, we would like to thank the following people (presented in alphabetic order) for their important contribution to this document:

Kostas Anagnostakis *Niometrics*
Elias Athanasopoulos *Columbia University*
Lorenzo Cavallaro *Royal Holloway University*
Markus Kammerstetter *TU Vienna*
Georgios Portokalidis *Columbia University*
Vassilis Prevelekis *AEGIS Research*
William Robertson *Northeastern University*
Jonathan Smith *University of Pennsylvania*
Edgar Weippl *Secure Business Austria*
Hossein Zakizadeh *Volvo*

Part I

Threats

1 Introduction

In the first year of the *SysSec* project, we defined three working groups, focusing on three distinct but strongly connected areas: *Malware and Fraud*, *Smart Environment*, and *Cyberattacks*. The purpose of each group was to discuss both current and future threats in the area of system security, and to contact and involve in the project a number of external experts.

The same structure was maintained during the second year, in which the working groups reached full maturity and worked closely together with the three research workpackages (WP5, WP6, and WP7). In this first part of the deliverable we present a summary of the evolution of the threat landscape, as seen from the three groups. The result is an updated version of what has been presented in D4.1: *"First Report on Threats on the Future Internet and Research Roadmap"*, maintaining a focus on the short- and mid-term future. Unsurprisingly, neither the research community nor the attacks developed by miscreants took a 180 degree turn during the past year. As a result, most of the threats discussed and presented in the previous document are still valid.

However, new topics also emerged during this second year. For example, the Malware and Fraud group added one new threat related to the possible spread of portable malware, i.e., malicious code that can run on multiple platforms. The Cyberattacks group largely discussed the problem of big data security, and the importance of the social engineering component in targeted attacks. In addition, the problem related to insider threats was again discussed with the help of a number of international experts in the field. Finally, the working group on Smart Environment focused this year its analysis on the areas of connected cars and smart grid. The problems that are identified are similar to the ones already described in the first deliverable (D4.1), but these environments also present some unique characteristics that affect the possible countermeasures.

Another interesting aspect of this second threats report is the increasing overlap between the discussed topics. For example, mobile phone security is an issue both from a malware point of view and from a cyberattacks perspective, and it is in fact mentioned both in Chapter 2 and in Chapter 4. Another point in common is the hardware security, that links together the research in smart environment and malware and fraud. We believe that the fact that similar threats are identified independently by researchers working in different domains is a strong support in favor of the importance and relevance of such threats.

The next three chapters present additional details about the sources of information adopted by each working group, as well as the output of the updated list of threats.

2

Current and Emerging Threats in Malware and Fraud

Contents

Analogous to the first Research Roadmap presented last year (D4.1), this document continues the process of assessing the current situation where Malware and Fraud is concerned. The input originates from discussions among the SysSec Project members as well as the second working group meeting, held on the 5th of June in Vienna.

2.1 Background

Naturally, one of the most important questions when reading this deliverable is how accurate the threats discovered last year are compared to the reality. Answering this question and explaining how the research community is aligned to these threats is the main purpose of this section, and of the deliverable as a whole. Most of the threats and research directions presented in the area of malware and fraud are still valid. There are, however, certain trends that take shape on both sides of the spectrum and are strongly correlated to the roadmap we presented. In the following, we discuss most of the threats introduced in the original document together with an analysis of its current status and activities.

2.2 Threats

2.2.1 Mobile Malware

The most blatant activity for both researchers and attackers has happened on mobile devices. Unlike desktop computers, powerful mobile devices which are capable of running full-blown applications are relatively new. Around 2006, smartphones became powerful enough to become interesting for malware writers. Before a new platform can be exploited on a large scale, however, the user base has to be large enough to gain the interest of malware authors (see Section 2.2.4 for more details). That may be one reason to explain why large-scale epidemics are yet to be encountered. On the other hand, most of today's mobile devices are either running Apple's iOS or Google's Android, with Android holding the majority in 2012 [15]. With over 100 million smartphone subscribers in 2012 in the US alone, the required user base certainly exists. However, the operating systems running these apps come with two major advantages when compared to PC operating systems:

1. They were developed in a time where security was already seen as important. Therefore, each major platform, be it Andropid, iOS or Windows Phone, come with a more or less advanced security stack and features to prevent these devices from easily being exploited.

2. Even more important is the fact that, currently, all mobile devices use a market to provide applications for the user. Since these markets are supervised and applications rated by users, it is harder for an attacker to distribute malware over these channels.

As a result, virus-infested programs are rarely found on official markets [6]. Instead, most samples stem from other sources like alternative markets (Cydia, Aptoide, etc..), torrents, or even direct download sites like rapidshare or similar pages. Furthermore, the relatively closed environment of these devices reduces the attack surface of drive-by-downloads on smartphones to an insignificant level. Still, the huge amount of potential victims is a reason good enough to investigate for miscreants. In fact, a rise in mobile malware is evident. Where traditional attack surfaces (spam, drive-by-downloads and remote exploits) are infeasible, new methods are developed [2]. Some examples are:

- Native code libraries to circumvent auditing features and hide from sandbox tools.
- Trojans in illegal games/apps, distributed by alternative markets.
- Specific exploits for rooted/jailbroken devices to circumvent security features the default device would have.
- Hidden functionality that exploits the fact that users tend to ignore permission dialogues.

These are just some examples of possibilities to still cause harm to smartphone users. Not all of them are already widely used but we predict that they will be in the coming year. One threat, which was specifically mentioned in the previous roadmap dealt with mobile trojans capable of intercepting transaction codes sent by netbanking applications via SMS. In February 2012, the first version of ZitMo (Zeus-in-the-Mobile) was encountered. The sample essentially behaves as predicted, sending SMS information (i.e. mTAN numbers) to a centralized service.
For the future, we expect more sophisticated forms of mobile malware. Depending on the success of the upcoming Windows 8 and its mobile counterpart Windows Mobile 8, we even expect multi-purpose malware apps that leverage functionality from both platforms, since they operate on the same basic kernel.

2.2.2 Malicious Hardware

The second threat we discussed in the previous roadmap deals with malicious hardware. Initially, the term was meant to incorporate embedded

devices, ASICS or even complete CPU's with dormant branches of functionality. Several ways to introduce these malicious circuits into an otherwise functional chip were discussed. They included possibilities like creating them in third-party components or even introduce them for specific, targeted attacks. While these possibilities certainly still exist, it is usually hard to prove their existence in a fabricated chip [10] which might even be protected against reverse-engineering. Furthermore, an investigating party would have to invest a tremendous amount of resources, both manpower and financial, to successfully conduct such tests. As a result, hidden functionality in hardware modules might already be there, but discovering and proving it is a different problem. Still, this form of hardware modification still poses a threat.

2.2.2.1 Test facilities

Contrary to ill-intended circuitry, recent research [13] has revealed an astonishing new perspective in hardware security. All companies that produce hardware chips utilize sophisticated testing procedures to ensure their product's quality. To ensure that each circuit is tested properly, on-chip testing facilities are utilized and the testing procedure itself is fully automated. After testing, these testing facilities are deactivated, by un-soldering JTAG connectors, opening connectivity fuses or flashing on-board memory with the respective directives. Unfortunately, reactivating these access points is possible in some cases, resulting in devastating effects for the security of the chips. In [13] for example, the authors re-enabled the on-board JTAG functionality for a military device, effectively enabling them to unrestrictedly access every functionality on it, including a complete reprogramming of the internal firmware.

2.2.2.2 Hardware bugs

Other hardware-level properties are also known to cause the underlying security stack to be compromised. In March 2012, for instance, car thieves exploited the ODB (on board diagnostics) port of a BMW to steal the car [12]. Apparently, the board was not designed to consider certain curious combinations of input parameters, which in turn, allowed the miscreants to access the vehicle.

Similar exploits are thinkable for all kinds of hardware implementations reaching from RFID-based access control to e-money. Even if the line between hardware and software blurs a little (is it circuitry or firmware?), these attacks will gain more popularity as technology penetrates other facets of daily used devices.

2.2.3 Cloud computing

Cloud computing still is one the most hyped IT innovations. Most IT companies announce plans for or already have IT products according to the cloud computing paradigm. The same is true for private users that decide to utilize a cloud-enabled service like dropbox, google drive or any other of the uncountable new instances that popped up around the IT landscape. It is already evident that its most critical flaw, according to public consent [8, 3], is security. From the predicted threats presented in the previous roadmap, API-level attacks were those with the least effort required for an attacker. The prediction proved to be true and, indeed, certain attacks were seen that targeted cloud services on API level [11]. Fortunately, these kinds of attacks are not among the very severe. In most cases, fixing the underlying API mitigates the threat or even completely removes it. Furthermore, API-level attacks are always valid for one service only. Therefore, a successful attack does not necessarily mean that other cloud services are compromised.

A scenario which has not yet been seen in the wild are attacks against the hypervisor to target cloud computers. The reason for that is that not every cloud uses the same virtualization technology. Therefore, an attack had to be targeted at a specific service again. And even if an attack could be found that breaks virtual machine boundaries, it is no guarantee to have access to all data from the rest of the service. Even for the prerequisite, which is virtual machine jailbreaking, no widely used attack vectors exist currently.

Finally, the most immediate threat is still represented by ordinary attack vectors that can be leveraged by cloud computing. Many companies either switch to cloud services for various purposes (data storage, pervasive computing, etc.) or implement their own solution. As a result, a compromising attack often leads to devastating impact on the structure as a whole. Instead of a single machine, the whole service is compromised, threatening all participating users at the same time. Another drawback of cloud solutions is the problem to recover to a saved state. If the underlying system is flawed, it can lead to severe downtimes for a company.

2.2.4 Portable Malware

A new threat, which was identified by the working group on malware and fraud is portable malware. One thing that stayed pretty stable in the years 1995 to 2010 was the operating system. Windows with all its flavors, was by far the most widely used OS worldwide. Naturally, malware also evolved for the largest target. Only recently, other platforms (Android, iOS, MacOS, or even Linux) are used by enough people to provide an interesting attack surface. Other than the user base, most of these platforms have an additional advantage for malware authors. The same kernel is used among different

target devices. iOS, for instance, is designed for both iPhone and iPad alike. Some even work on MacOS. Therefore, also malware written for these platforms reaches a broader mass. The same is true for Android, which runs on mobile phones, tablets and even laptops. The only advantage so far is that a relatively low number of samples exists for these platforms and that the attack vectors are smaller (see Section 2.2.1). With the upcoming Windows 8 and its mobile pendant, Windows mobile 8, however, we expect a lot more portable malware samples to reach the broader community.

3

Current and Emerging Threats in Smart Environments

Contents

The second meeting of the Smart Environment expert group took place in Vienna, June 2012, to discuss threats related to the area in question. The deliverable, "First Report on Threats on the Future Internet and Research Roadmap" (2011) and a draft of the deliverable "Securing the Connected Car" were used as a basis for the discussion, where the objective for this year's meeting was to consider the continued relevance of the threats identified the previous year and to discuss two key areas in more detail: the connected car and the smart grid. Overall, the experts found that the threats identified last year still remain valid. It was noted that the scenario, "The peccadillo", is highly relevant. This particular scenario describes how a consumer hacks her own smart meter by installing a custom firmware to reduce her electricity cost. In this particular scenario there is an unknown side effect of the firmware that allows an external criminal organization to control the smart meter. A cyber intelligence bulletin from FBI obtained by KrebsOnSecurity in April 2012 describes how smart meters have been hacked to reduce the energy consumption of the customer, resulting in a large financial loss for the energy company.[1] In this particular case, it is suspected that the optical port of the smart meter was used for the reprogramming of the device. According to the bulletin, only a moderate level of computer knowledge is necessary to compromise meters and the tools and software required do not cost much and are readily available on the Internet.[2] With the large-scale deployment of smart meters that is planned for the next decade, it is likely that attacks targeting this environment will increase.

In the following, we present the background and the context for the smart environment expert group. We then shortly summarize the findings from the previous meeting. This year, we discussed *the connected car* and *the smart grid* in detail, not only focusing on threats per se but on general problems and possible mitigation techniques that could work for these areas. The discussion is then summarized and the chapter concluded.

3.1 Background

The focus of the smart environment expert group is low-capability devices, ranging from simple sensor networks to more heterogeneous systems with more capable hardware. As there is a continuous range of such devices and what they are capable of, a threat and the corresponding mitigating security mechanism may look very different depending on the type of device and the environment it is located within.

In the discussion last year, the general characteristics of typical smart environments were discussed, especially the challenges in SCADA systems, the smart grid, and the connected car. For example, what are the trends

[1] http://krebsonsecurity.com/2012/04/fbi-smart-meter-hacks-likely-to-spread/
[2] http://www.blackhat.com/usa/bh-us-12-briefings.html#Weber

concerning the capabilities of these devices? It is expected that some devices will increase their capabilities in the future, but, as pointed out by the experts, certain parameters will not change much over the next couple of years. Even though new nodes will run on better hardware, using less power, power management will remain of paramount importance for sensor networks. Even though a few more bits may be used for encryption in such environments in the future, the fundamental properties will not change much and the security solutions need to be adapted to the special requirements of the environment in question, meaning that power management will still have a major influence on every piece of code running on nodes.

It was highlighted, both in the discussions last year and in the discussions in Vienna this year, that it is expected that the use of devices in smart environments will increase over the coming years. These kinds of devices will also be found in areas where their correct function is of utmost importance, such as in critical infrastructures in society. For that reason, security needs to be emphasized, both for new deployments with hardware and devices that will have a long life cycle (up to forty years in some environments), but also for legacy deployments. Proprietary and closed solutions should be avoided for the same arguments as with custom-built encryption algorithms and open, well-researched solutions that are standardized should be the norm.

3.2 Threats and problems

Below is a summary of some of the characteristics regarding smart environments that may pose a threat. These issues were discussed in-depth last year.

The *accessibility* of the actual devices and the networks, be it either physical or logical, may be a problem. The smart meters are located on the premises of the customers and, as described above, the customers may tamper with the devices for their own benefit.[3] In a sensor network, many nodes are in exposed or accessible areas. For the case of the connected car, either the owner or even someone close to the car may attack it. Researchers have already shown that it is possible to attack vehicles, for example via the multimedia system in the car (or the example given in Section 2.2.2.2 concerning the theft of a BMW). This problem will become even more serious when vehicles begin to communicate with the outside world.

The *system complexity* is also significant for typical smart environments as they consist of not-so-capable but numerous devices. These devices are very cheap, meaning that the cost of security per device must be low and there is no user interface to many of the devices and no central control or management point.

[3] http://krebsonsecurity.com/2012/04/fbi-smart-meter-hacks-likely-to-spread/

Systems comprising smart environments are many times costly to update, making *maintainability* an issue. Sensor nodes may be located in inaccessible areas which are impossible to get to, or the scale of the deployment may make any type of hardware change on the individual nodes impractical (hundreds of thousands of smart meters in a city). Some systems allow wireless remote reprogramming to maximize flexibility but such a feature can in turn be used for attacks.

A slightly non-intuitive vector for new attacks may be that *more capable devices* will be deployed in the near future. From a first viewpoint, this should be positive as devices with more cores, using less battery, may utilize better security primitives. However, the experts see an associated risk in that higher level programming will be used to speed development of software for the devices. The carefully handwritten code of today will be replaced with a simplified java or even a scripting language bringing with it new sets of vulnerabilities as seen for regular ICT systems. A specific case that was discussed separately is that many consumer devices get more capabilities to communicate over different frequencies using different protocols. Given that there are cheap and *ubiquitous readers* available, a larger scale of fraud and other types of attacks may result that before required specialized hardware. A specific example is the development of near field communication readers and emitters on cell phones that may be used to falsify tags.

The *network layer protocols and services* need to be secure. It is important to have secure algorithms for all the basic services that are needed in sensor networks, such as: i) Routing Protocols, ii) Aggregation, iii) Localization, iv) Clock Synchronization, v) Clustering, and vi) Key Management. However, in some smart environments, legacy protocols are used to increase backward compatibility, or proprietary protocols developed by a single vendor is used with no guarantees for properties important for security.

Finally, cyber physical systems are comprised of both an ICT component as well as other sensors or actuators that control the environment. In the future, *attacks against the non-ICT component*, or attacks that target both the ICT component and sensors simultaneously may occur.

3.3 The connected car and the smart grid environment

3.3.1 Special characteristics

The first question posed to the expert group this year was if we are simply solving the "same" type of problem over again? Cannot the mechanisms found for more traditional systems also be used in smart environments? Last year, it was argued that many traditional security mechanisms do not work for systems in the smart environment, either because the underlying

assumptions are not valid (no patching possible) or because of more practical reasons (proprietary protocols). In the discussion this year, the experts overall agreed that many smart environments have special properties that make simple adaptation of known techniques difficult. The following special needs for the smart environment were emphasized, thus expanding the properties identified last year. It is important to run on low power with a minimum of communication overhead. Cost is often an issue, permeating the design from start to finish. Moreover, some of the areas where smart devices are found are also governed by compliance laws – the metering device for the smart meter needs to fulfill a number of criteria set by the government in a country if it is going to be allowed to measure the electrical usage of customers in that country. In Austria, only a certain number of brands have been certified. Even though such compliance criteria overall are beneficial, they also add overhead to the process. The certification of devices may be costly and thus, changes (bug fixes) may be introduced at a slower rate than in other environments that are not as strictly controlled. Another example where laws may govern the technical implementation is with car to car communication. In most cases, this communication should be anonymous but there may be cases in certain jurisdictions where the messages need to be released after a court order.

However, the experts also pointed out that it is important not to generalize between the environments as they range in their capabilities. For example, even though power may be an issue it is important to realize to what degree and when. For a smart meter, connected to the grid, power is available unless there is a blackout. For a car on the highway, there should be enough power to drive a more secure protocol but care should be taken if the engine is turned off. For a cell phone, the most invasive memory scanning techniques should only be performed when the phone is recharging. The energy to drive a more secure protocol or scanning technique will always come at a cost and, today unfortunately, many companies or consumers may not be willing to pay for security.

3.3.2 Black box implementation

The discussion of the different types of environments brought the next problematic points for the complex commercial systems that are deployed today. These systems are often proprietary and from an external observer (researcher) the implementation is a black box. It is difficult to find out any details of the hardware or the firmware of the devices, making the development of mitigation techniques challenging. The protocols are also many times proprietary or built on an older protocol with special vendor extensions. Thus, there exists very little open information of different systems. Each vendor has produced their own system.

3.3.3 Missing toolset for analysis

Moreover, due to the black box implementation of the system as well as its proprietary nature, there exist no tools for investigating such systems. For PCs or smart phones, there are toolsets available to researchers and a clear methodology that can be followed. SCADA systems or smart meters are very difficult to analyze in comparison. Many of these systems are created by companies with little or no known experience in ICT security; they are experts in building control systems but the question is whether they know enough of security to build robust devices that can withstand different types of attacks? The requirements of the systems have changed, where the devices need to communicate over IP with the proprietary protocols on the application level, and thus there is a much higher risk that the system will be attacked. Security is complex and sometimes laymen have the attitude that as long as they use a standardized cryptographic protocol, their product will be secure. As independent researchers lack the proper information and toolset to analyze the system, the public is left with just trusting that the companies have done the right thing and have built robust devices that can withstand different types of attacks.

The question was also raised how vulnerabilities in these systems should be handled. Certain security weaknesses may only be fixed by changing the hardware or manually updating the firmware. However, given the planned large-scale deployment for, for example, the smart grid, such operations come at a substantial cost and will take time.

3.3.4 Possibility of open-source components

The analogy to the development of cryptographic algorithms was actually seen as one way to build better systems within this domain. Cryptographic algorithms today are mostly developed through an open process and closed designs are not promoted. As the devices found in the smart environment are many times going to be used in critical infrastructures, it is vital that a similar process is used. It is important that the "standard components" such as the network stack is open source and available to all researchers for inspection to make it as secure as possible. With standardizations, it is possible to build several pieces in an open source manner that vendors can use when they implement their own system.

As a comparison, the relatively open environment of the car was discussed. The car may be a slightly more open area compared to SCADA systems and other systems found in the smart grid. With the car, there are numerous standards that need to be followed, allowing some interoperability. However, there is still the problem of backward compatibility of older components as the life cycle is quite long. The open environment in itself also creates new challenges. Parts of a car may be sold by the car company

itself but also by a number of authorized vendors and installment can be done by the customer herself. For smart meters, the customer does (so far) have no choice in the matter but the energy company decides the type of smart meter that will be installed and the functions that are allowed. As a parenthesis, there exist other systems where open access is both necessary and important, such as medical devices. The monitoring of an insulin pump should be straight-forward in an emergency room.

3.3.5 Drivers to make the environment more secure

So far no repeated large scale attacks have been seen against these types of environments. However, if the attack is large scale it is probably visible and will be noticed immediately. If the attack is targeted, it may not be noticed and can lie dormant until certain triggers occur. If the latter has happened is difficult to tell at this point. Vulnerabilities have been documented but the question is if they have been used by nefarious groups.

This session concluded with a discussion of the necessary drivers to make such environments more secure. Standardization is important, government policies and compliance to a set of laws will influence a minimum set of security requirements. But equally important is the availability of an open source tool set and knowledge so that researchers can properly analyze the systems in question. With large scale deployments in critical infrastructures, the system and devices need to have security by design that can be tested by independent researchers. The public cannot simply blindly trust a commercial company and assume that they have managed to solve all security problems that can occur in these environments.

3.4 Summary

The objectives of the second expert meeting were twofold. First, were the general characteristics and the resulting threats identified last year still valid? Not surprisingly, the experts agreed with the list as these threats are not likely to change on a year by year basis. Second, the experts then looked more closely at the environments of the connected car and the smart grid to identify further issues. It was noted that a number of systems are of a proprietary nature and sometimes the systems can only be seen as a "black box" from the point of an independent researcher. There are neither detailed information of the system or a toolset to analyze it, meaning that any analysis is difficult and time consuming compared to more traditional ICT environments. This implies that we might not know enough about the weaknesses of the systems that control critical functions in society. As some systems are already deployed on a large scale, the question is how weaknesses identified in the future should be handled. The normal patch cycle of

more traditional ICT systems may not be applicable as actual hardware may have to be replaced to a great cost. However, attacks during the last couple of years have highlighted the need to study the security of the systems we use in our critical infrastructures.

4

Current and Emerging Threats in Cyberattacks

Contents

The threats presented in this chapter are the results of brainstorming and discussion within the *SysSec* Project. Specifically, within the *SysSec* Working Groups at the second *SysSec* Working Group meeting in Vienna on the 5th of June 2012, as well as discussions with members of the Cyberattacks working group and other experts in the area of Cybersecurity. The following section presents a revision of the respective section in Deliverable D4.1, "First Report on Threats on the Future Internet and Research Roadmap" (2011). We believe, the threats discussed in the remainder of this section will have increasing impact in terms of security in computing systems and networks in the following years.

4.1 Background

The focus of the Cyberattacks working group is to improve our understanding in new and emerging types of cyberattacks, such as attacks on and by mobile phones and other such highly-connected smart appliances, web attacks, attacks on home and office automation devices, cross-domain attacks, attacks on individual citizens as well as infrastructure, etc. It is also the goal of the working group to advance the State-of-the-Art in the area of detection and mitigation of such cyberattacks.

4.2 Threats

4.2.1 Social Engineering

A new topic that came up with the discussions with the experts is social engineering. Obviously this is a well known and widely accepted threat. The main reason this was not included in the previous report is the acceptance that even though it is a very serious and important problem, it is hard to address with purely technological means.

The topic was revisited and discussed again by the experts and the partners of the project. The discussion revolved around the idea of using capabilities to assist in the defense against social engineering attacks. Even though they are hard to program, they may prove to be effective to a degree, as finer grain control over resources may be hard to get around even if one successfully manipulates a user into performing certain actions.

4.2.2 Web Services and Applications

The second topic of discussion was that of the proliferation of new web services and the plethora of new applications. These new services and applications will continue to attract new users, and as this happens they will inevitably become targets for the attackers. Due to the quick evolution of

this new model for deploying software we expect that they will remain vulnerable due to software bugs, misconfigurations and the desire of users to install and use them.

This of course will have direct consequences on end user security as we have come to depend on these online services in our daily lives. For example, by compromising a news service, miscreants may spread misinformation which can have direct financial and social impact. By taking down government web services relating to tax or other internal revenue, one will cause major impact on a country's economy. By infecting an online storage service, individuals or organizations may lose important data stored online.

Also, as already discussed in Chapter 2.2.1, the inverse types of attacks are possible. That is, miscreants can use infected or otherwise compromised online services to attack all types of end appliances. This has been traditionally possible against personal computers, but as users start using their phones, tablets or other smart devices to synchronize their data with online sources, or download applications for personal use, we expect these types of attacks to increase.

4.2.3 Big Data and Privacy

The next topic that the Cyberattacks working group focused its attention on is that of data collection, data aggregation, data usage, and the effects it has on citizen privacy. The problem becomes apparent if one observes that Internet is an invaluable source of information about every conceivable topic. However, in recent years, data put on the Internet have evolved from purely encyclopedic information about a variety of topics, and simple user pages, to much more personal information. This trend has been facilitated by the growth of social networking sites. A social network is a social structure that is made up of nodes that represent individuals or organizations. These nodes may be tied to each other by properties such as friendship and general interests. As these online communities, such as Facebook, MySpace, Orkut, Twitter, LinkedIn, and others, have been adopted by Internet users, miscreants have started abusing them for a variety of purposes, including stalking, identity theft, spamming, direct advertising, spreading of malware, etc.

The reason such attacks are possible, is due to the nature of information users upload to social networking sites. Users typically give their e-mail address, where they went to school, what they studied, jobs they held, places they lived, their relationship status, family information, their friends, hobbies, places they have visited, likes and dislikes, etc. There is really no limit to the amount and detail of personal information users will upload. From the attackers' perspective this is fertile ground for learning about their victims. The e-mail addresses can be used for spamming, friend information can be used for targeted attacks, and data about other habits can be used for blackmailing.

The attackers can also correlate information from multiple social networking sites, along with other sites, such as blogs and online forums, and even documents such as Word or PDF files, to really learn things about their potential victims [4, 5]. The more information they hold, the more likely it becomes that they can somehow exploit their target.

There are a number of challenges here. How can we protect users from the constant leak of, possibly private, information. Do we have the right metrics to quantify the threats and risks from the accumulation of so much data? Are there ways to implement "the right to be forgotten"? How are new technologies as the ones that are deployed with the connected car, and smart environments in general going to affect personal data collection and privacy. The work group experts remained convinced this will be a growing and important problem in the following years.

4.2.4 Critical Infrastructures

The border between what we traditionally considered critical infrastructures and the public Internet is quickly disappearing. Change is taking place in both directions. That is, on one hand, critical infrastructures are becoming more connected to the public network, on the other, ICT infrastructures are becoming ever more necessary to our daily lives.

For example, one can think of the telephony network as a traditional critical infrastructure, used by billions to communicate. However, what we are witnessing is an ongoing migration towards VoIP services, effectively eliminating the line between the telephony network and the data network. Recent work has shown how one can exploit VoIP services to attack emergency service land lines [9].

The same applies for other technologies as well. For example we can consider the case of data centers, and cloud computing infrastructures in general (see Chapter 2.2.3 for a discussion on this topic). Such environments host numerous services used by thousands of businesses and millions of users. This makes them ideal targets for attackers. Taking down a cloud provider, or penetrating their infrastructure and stealing or modifying data, can lead to serious disruptions, and possibly millions of Euros of damages. Currently we are not trained to view or consider these online services as critical infrastructure, in the same sense as we view the electric power grid as critical infrastructure.

The working group discussion around critical infrastructures moved beyond what has been discussed so far into the area of the connected car. There was active interaction between the Cyberattacks working group and the Smart Enviroments working group. There is extensive discussion on this in the previous section. What is important to note is that the set of what is considered critical infrastructure continues to grow. The threats and risks posed will continue to be of importance. The discussion also was steered

towards fuzzers and how they can be applied to systems in order to increase their robustness. We believe there may be avenues there to explore, much in the way as fuzzers have been used to improve application software by exposing its weaknesses.

4.2.5 Smart, Mobile and Ubiquitous Appliances

The adoption of all types of smart and mobile devices may be posing a more serious threat to security than any other time in the past. In the United States alone smartphones are vastly outselling traditional PCs.[1] This makes past scenarios about devices and sensors, static and mobile, being deployed universally, a reality. Obviously, such devices have varying characteristics, but their underlying common features are: ever-increasing computational capabilities and continuous connectivity, be it Ethernet, WiFi, GSM, 3G, Bluetooth, radio, or even infrared.

These devices take many forms, that may rarely remind us of the traditional personal computers we are so used to, but in reality they are very much vulnerable to similar types of attack vectors, customized to each specific device. For example, medical appliances such as pacemakers, have been shown to be vulnerable to attacks [7]. Such vulnerabilities may lead to direct loss of life.

Attacks however do not need to be directly threatening to human life to be serious in nature. Smartphones are a case in point. Nowadays, our phones hold a treasure of sensitive information: phone numbers of our family, friends and colleagues, personal photos, financial data, passwords, virtual cash, location information, etc. In some respect, our phones may be a more valuable target to attackers than our personal computers or servers.

Malware taking over our phones, we believe, is a very real threat. Malicious applications that users install without realizing their true intentions are one of the possible sources. As users are willing to download and run programs from online sources on their smartphones, they become trained to accept without thinking pretty much any request the application may make. For example, access to the network, to storage, or even debug mode of the phone. This leaves users vulnerable to software that may provide some surface functionality, e.g. a game, and stealthily steal information in the background.

Exploiting such devices is often easy due to a number of factors, not all applicable in all cases: limited computational power to run full fledged security software like antivirus, firewalls, or intrusion detection systems, dependency on battery power, so even if security software exists it may not be practical to run, lacking security design, ease-of-use trumping security

[1] Ars Techinica: "From Altair to iPad: 35 years of personal computer market share", http://arstechnica.com/business/2012/08/from-altair-to-ipad-35-years-of-personal-computer-market-share/4/

requirements, easy physical access by attackers, etc. These characteristics have attracted a great deal of attention by the research community. In the past year, and in the years leading up to today, security researchers are allocating more effort into research about security mechanisms for smart devices.

4.2.6 Insiders

Another topic that the expert group discussed vividly was that of the malicious insider [1]. This is an often overlooked factor in cyberattacks. Insiders such as: opportunists, disgruntled employees or even malicious plants from competitors and adversaries, all pose tremendous challenges for ICT security. Typically, organizations follow the model of forming a strong perimeter to repel attacks coming from the outside [14]. This is expected, as traditionally insiders are considered trusted by the mere attribute of already being on the inside. Unfortunately this is not always the case. Employees change position and move from one department to the other, new ones are hired, some leave and never get their privileges revoked.

Insider attacks are more dangerous than attacks from outsiders, as insiders probably have easier and more direct access to the assets they aim to compromise. Additionally, they may already know of the countermeasures put in place, or have other intelligence that will help them in their goals. Furthermore, security mechanisms are typically tailored to counter outsiders. These are placed at choke-points along the perimeter of an organization. Once inside, very little defenses are in place. To make matters worse, insiders also have a lower chance of getting caught, since as we said, defenses are along the perimeter, but also because we are trained to look to the outside for malicious activities.

Once an insider goes rogue, they may sabotage the organization, for example by modifying or deleting data, locking out computing systems and networks, etc. In these cases, the malicious insider may be easier to detect and track. In other cases, where the malicious insider has more long-term goals, they may start stealing the organizations' intellectual property. Such attacks are harder to detect, and even if detected, an organization may not be willing to admit such events.

Due to the above, it is imperative for organizations to form policies and implement controls that monitor, detect and prevent access to sensitive resources, irrespectively of who may be considered trusted or not. In the discussion with the experts several points were brought up. First of all, security agencies have procedures that they need to follow and deal with malicious insiders. Sometime background investigations can be used to weed out insiders. Banks deal with insider threats by forcing people to take two weeks of vacation consecutively. During these two weeks whatever the insider has

done will be exposed. All in all, the experts agreed that this is a topic that is worth investigating further.

4.2.7 Network Core Attacks: Here to Stay

There was general agreement in the discussion with the experts, that attacks against the network core will continue to exist. Attackers still find the core as an attractive target, the reason being, it is the glue that holds everything together. Compromising the core, or part of the core, may lead to other attacks. In a way it can be the enabler of attacks against the actual target of miscreants. Also, it is relatively simple to carry out a distributed denial of service attack against part of the core, which will in turn affect a large subset of the network.

Also, the fact that the Internet is the de facto unifier of a large number of communication services, makes the Internet core a valuable target. For example, the traditional telephony network is migrating onto the Internet, but also other services as well. For example, gaming has been steadily moving online for a number of years now. There are other forms of entertainment that are moving online as well. For example, television shows and motion pictures are appearing online today more than ever before in the past. Such services are prime candidates for today's attackers, and a simple way of attacking such services is by taking down the underlying functionality inside the network core.

The consensus of the discussion group was that we will keep on seeing attacks against the network core. They may not be as impressive as in the past, as attackers start to focus against specific web services and applications, and end devices such as smartphones and tablet PCs.

4.3 Summary

The objectives of the Cyberattacks working group was threefold. First we wanted to go over the list of threats from last year's report, and evaluate which of the threats are still relevant in the new landscape. Secondly, we wanted to update the list of threats with possibly new ones that the group believes will be appearing, or reappearing, on the horizon. Lastly, we wanted to interact more closely with the experts of the other working groups and see how the threats overlap between the three areas.

As an outcome of our work, we want to stress that there haven't been any major changes in the threats landscape. We believe that the threats from the previous report are still relevant and important today. Even more so in cases where there is closer integration between open networks and critical services, where there is increased adoption of new applications, and when the bulk of data about citizens and users of the Internet is increasing. To our

list of threats we decided to add that of social engineering. We think this is a hard problem to solve, and will remain relevant and be used by miscreants, when other means of attack, such as exploiting specific vulnerabilities, is not an option.

Part II

Scenarios and Research Roadmap

5 Overview

In this second part of the deliverable we present three new attack scenarios, based on the threats discussed in the previous chapters. The ones we presented during the first year of the *SysSec* project became extremely relevant in 2012. As it was already mentioned in Chapter 3, the "peccadillo" scenario, in which we described how a consumer could hack her own smart meter by installing a custom firmware to reduce her electricity cost, almost perfectly matched a cyber intelligence bulletin from FBI obtained by KrebsOnSecurity in April 2012, describing how smart meters have been hacked to reduce the energy consumption of the customers. The diffusion of the new version of Zeus-in-the-Mobile (see Chapter 2 for more information) was instead tightly related to the second scenario we proposed: the bank job.

Again, what we present in the rest of this chapter is not something that could happen in a remote future, but instead something that could potentially be observed in the wild either nowadays or in few months from now.

Finally, in the last part of the deliverable, we present the new approach we adopted to update and define the research roadmap. This year, we followed a more formal methodology, introducing a new graphical representation to summarize several characteristics of each direction (including a measure of the impact, likelihood, time frame, technological challenges, and more). The SysSec consortium confirmed the importance and the significance of the first roadmap, but each research topic has been updated to reflect the small changes in the threat landscape presented in this document.

6

Scenario: The Contact Dealer

Contents

6.1 The Story

Social networks (SN) became extremely popular during the last years. They contain hundreds of millions of active accounts. An important difference between SN and earlier forms of social interaction such as anonymous or semi-anonymous newsgroups and forums is that the people tend to share much more true and up to date information about their personal or professional life. In short: the people are more likely to be honest in the current social networks. Moreover the SN are open and public, and everybody can easily register. On the other hand, the corporate networks are often restricted, and sometimes even hidden. Actually the SN may contain significant personal information which is not stored even in the confidential records of a corporate or government network. Therefore, SN are a very good target for social engineering. This story story presents a general scenario that may affect every user, and it is not focused on a social network in particular.

John, the attacker, is a computer hobbyist with moderate IT skills and knowledge. He owns three computers, a laptop, a smart phone, and some other useful and cheap hardware such as a printer. Software needed: nothing very special above the basic distributions of MS Windows and Linux, as well as a reliable Internet connection. But above all, he has plenty of free time and patience! And he possesses a wide general knowledge, eloquence, a rich vocabulary and some psychological talent. John knows the features and operation of the popular SNs and understands that the information contained in them is very valuable and has previously been unavailable. The main purpose is to gather the maximal amount of personal information and to gain the trust of several real persons that have important or interesting positions.

6.2 Explanation

The first step that John makes is to create a profile on some of the most popular social networks. This profile must be comprehensive and convincing. It is not a good idea to present himself as an IT specialist of any kind. It is possible to create another different profile, for actions of type “split personality”, but it is dangerous because it requires far more attention and a good memory during the communication with the SN users. A successful usage of two or more different personalities in a SN is a rare case probably. Usually one well defined profile is enough. Next, this profile needs to mature. Maybe the first thing that a SN user looks in a profile is “Registered on: ...” (if it is provided by the SN). The very new profiles are at least a little suspicious for a period of time. Therefore John spends about three months in light and enjoyable activities in the social network, so that his profile gets old enough.

It is time for purposeful actions. John sends friendship requests to several people. Let's say 100 are enough for a start. He tries to win their trust over the next three months using several approaches including the ones already mentioned above such as eloquence, wide general knowledge and the most important – patience. The hardest part of the work is done when a majority of these initial 100 people start to trust John. If John finds out that the target is not "interesting" he can drop it and replace it with another profile. The idea is to get in contact with several people that are employed in critical sectors and companies (e.g., medical, universities, government offices, or just any company that has valuable assets).

Now follows a more technical part of the exercise. John needs several software tools which will help him to collect, store, verify and sift out the information. Also he keeps some kind of diary to track his activities in order to avoid contradictions during the conversations with the victims. Fortunately, most of these tools are not complex and are already available in the basic OS distributions - simple RDBMS, script and text processing such as shell scripting, PERL or PHP, other built-in programming languages, editors. The design of the RDBMS is not strictly specified, it contains many fields which are supposed to be interesting and it can be expanded during the progress of the attack if necessary. Each of these fields have the corresponding attributes 'true', 'false', 'unknown' to be used in the verification process. John develops some simple additional scripts using basic data mining principles. For example, one of these scripts is a crontab job which frequently, say every 15 minutes, monitors particular activities of a SN user. The script may send a fast notification, an SMS to John, if the user did something online. This is useful for speeding up the conversation and shorten the initial investigation time. John also collects and stores information for users which information is outside of the given SN. This includes accounts and nicknames for other online facilities, such as participation in specialized discussion forums, mailing list subscriptions, RSS feeds, own and favorite blogs, personal webpages, instant messengers, torrent trackers, used smart devices etc. These facts are revealed during the chats and using the regular search engines and cross-checked with other sources, when it is possible. In addition, John can also gently and without insistence, invites some of the friends in face-to-face meetings. This approach is not easy to conduct, especially when the users are not concentrated in a relatively small geographical area. But the main advantage is that many more facts can be discovered and verified in such meetings.

After John collected enough verified and valuable data for hundreds of SN users. Now the main question is: How to cash the situation? Many individuals and institutions are willing to buy such kind of information – employee headhunters, marketing managers and analysts, and spam producers just to name a few possibilities.

But more importantly, John can sell access to certain profiles to criminals, who are interested in targeting that person and install malware on her computer. For example, many employees have VPN access to their company networks and, therefore, can be interesting victims of targeted attacks.

6.3 Final remarks

The described scenario is based on simple and well known social engineering techniques and psychological approaches. It works because the people tend to advertise themselves as much as possible in today's ferocious business world. Moreover the SN users seek social interaction and friendship. The main advantage of this method is that it does not require a strong high-tech background from the attackers. A disadvantage is the relatively long time before the first results are achieved.

7

Scenario: Portable Device in Stepping-stone Attack Against a Secure Network

Contents

7.1 Prologue

Computer viruses and malware in general need a transmission medium in order to migrate from host to host. This is usually the network itself, so the transmission is direct. In cases, however, when a network connection does not exist or is not convenient (e.g. the presence of a firewall, or IDS) the attackers resort to using storage devices such as flash cards, CDROM etc. This technique goes back to the beginning of the history of viruses, before personal computers were networked, when viruses used floppy disks as their transmission medium.

Nowadays, the stuxnet virus and an earlier virus that contaminated many hosts in the Pentagon, used USB flash memory devices to breach the external network perimeter. In the case of the Pentagon breach, the USB memory devices were left in the parking lot assuming that anybody who found such a device would connect it to their computer in an attempt to locate the owner. Thus, there is an element of human engineering involved in the attack as well. Nevertheless, memory devices are passive devices, relying on a human to connect the device to a computer, and the computer to execute code stored on it.

The proliferation of smart, hand-held devices (e.g. smartphones, or tablets) vastly enhance the options available to an attacker, since the transmission device is now capable of running malware that continuously evaluates its environment attempting to find a way to infiltrate hosts inside the secure network.

7.2 Attack Scenario

In this attack scenario, we look at a compromised smartphone that has WiFi capability. The malware has completely taken over the operating system on the smartphone, so that even when the device is powered off, the processor is still running the malware, while, to an external observer, the device appears to be powered off.

The malware may go into active mode if a prearranged event or condition is detected. For example, one such condition may be when the smartphone is powered off on the assumption that, in a secure environment, cellphones and other communication devices will be expected to be powered off. Other triggers may be a timed alarm (with the approximate time that the smartphone is expected to be in the secure location), receipt of an SMS, or even a device planted outside the secure location marking the location (this device may pretend to be an innocuous WiFi hotspot with an SSID that is known to the malware).

Once triggered, the malware scans the environment to detect vulnerable devices. The attack medium in this case will be wireless networks such as WiFi or Bluetooth.

In our scenario, the malware detects a printer/scanner that has its WiFi connectivity enabled by default. This printer is connected to the wired Ethernet network, but its WiFi interface has been inadvertently left enabled. This is by far a very common occurrence, since infrastructure devices are usually installed by a hurried support person who does not have the time to explore all aspects of the device in order to secure it.

Back on the compromised smartphone, the malware has configured the WiFi interface to respond to the connection requests from the printer and has managed to connect with the printer firmware using Internet Protocols (IP). The malware has a list of vulnerabilities of popular printers and quickly matches the printer model and software version (which the printer itself provides when queried by the smartphone) to a known buffer overflow bug. Again the attacker is helped by the fact that embedded devices such as the printer in this example, seldom, if at all, receive upgrades to their software. Thus, they often run early versions of their firmware with known vulnerabilities.

The implication is that the malware on the smartphone does not need a zero-day vulnerability to attack the printer, but rather exploits a two year old buffer overflow to inject malicious code in the printer firmware. Eventually the malware itself will migrate to the printer and copy itself to the flash memory of the device, so that even if the printer is power cycled, the malware will survive.

Since the printer is connected to the wired Ethernet, the malware can scan the network for vulnerable hosts, in order to spread further. Alternatively, it may just sit in that printer and relay copies of any documents sent to that printer back to the smartphone. Once the smartphone is outside the secure perimeter its malware will send the copied documents to its handlers.

Even with restricted memory and bandwidth (e.g. if the devices are using Bluetooth rather than WiFi) the malware on the printer may scan the submitted documents for keywords so that it does not blindly send everything to the smartphone, but only documents that match its list of keywords.

Once compromised, the printer will carry out its task even if communications with the smartphone are lost (i.e. when the smartphone is taken outside the building). When the smartphone is brought back into the secure building, the printer will have a collection of intercepted data to upload. This time the smartphone will have an updated version of the printer malware. This version will be customized to the particular device (and its capabilities) so as to further exploit the established bridgehead into the secure network. It also contains a new virus specially configured to attack a workstation, detected by the printer malware in the first visit, which runs software with a known vulnerability. If the workstation is also compromised,

the attackers will have gained access to the files on the workstation and any network file servers to which the particular workstation has access.

As the infection progresses, the cost of the eventual clean up increases, giving the attackers yet another advantage. Even if their malware is detected, the facility will have to be shutdown in order to clean up all the infected machines. Even so, the printer may slip though the net, thus opening the way for a re-infection once the network is back on-line. Thus the loss in terms of wasted man-hours is added to the cost of the lost information.

8

Scenario: Password reuse and mobile applications

Contents

8.1 The Story

John has been using computers for a long time. He owns an old PC at home, but his main and heavy interaction with the Internet happens every day through his office desktop. Like many organizations, his company prohibits the employees from visiting social networks and other entertainment resources during work. Thus, John can hardly enjoy all these new web applications. However, recently John purchased a smartphone and he happily replaced his old cellphone with a full-featured device equipped with GPS, video camera and accelerometer, running the latest Android operating system. John now had the opportunity to utilize all these social networks, such as Foursquare, LinkedIn, Instagram, Twitter and Facebook through his new smartphone, even at work. Without loosing a minute, John downloaded these social applications and signed up for new accounts. John was very excited about all this. He could post his location to Facebook, while attending live the local soccer derby, he could check-in in his favorite coffee store using Foursquare, and thus get a price discount, and, all these, through a tiny device of the size of his palm. John was extremely surprised when one morning he saw a tweet, originating from his own account, supporting his rival soccer team. Fortunately, it was only a joke made by his close friend, Maria.

8.2 Explanation

This is how Maria managed to steal John's credentials for Twitter, in order to impersonate him and post a message in his feed. John uses all web applications practically only through his mobile device. This is the case for many teenagers nowadays and people with heavy daily working routine, who can hardly access such resources using the corporate network. Each web site, such as Facebook or LinkedIn, comes as a stand-alone application. John registers to these services through his mobile device. Typing in smartphones is still a non trivial task. For example, mixing letters with numbers and special punctuation symbols requires switching the main keyboard. Also, typing long words is error prone. Thus, John makes three major mistakes. First, John selects a short password to be able to enter it quickly using his mobile device. Second, since smartphones are not equipped, at least by default, with password managers, John uses slight variations of the same password to many different services. Third, John knows a bit about technology; he is using services over HTTPS (and, thus, he is sure his password cannot be leaked so he selects a weak one), and he registers only with popular services, which can be considered trusted. However, even in cases where HTTPS is in place and services are considered trusted, passwords can be leaked offline. External intruders or inside attackers may target the service, and not

the user in particular, by exposing all storage containing password hashes in public.

One morning a group of crackers manages to obtain the passwords of one large social service provider. Passwords are not stored in plain text, they are hashed. Unfortunately, they are not salted. Storing a password usually involves taking the plain text and feeding it to a cryptographic hash function (like SHA1 for example). Only the output is stored and not the plain text. It is computationally hard to reveal the original password just by accessing the hash output. However, an attacker can use a large dictionary with common words, create various combinations and check if any of these words, if hashed, match with the hash in consideration. This is why a salt is used for making this task much more difficult. Instead of just hashing the password, a random number is attached to the plain text (for example, this can be the length of the username), and then the cryptographic hash is computed. This makes identical passwords having totally different hash values.

The other morning, Maria read in the news that the passwords of a huge service provider were leaked. The file with the hashes was posted in paste.bin and the achievement was announced in Twitter. As soon as the news came to surface, the file was removed from paste.bin, but the damage had already been done. Many users had downloaded the file and reposted it in hundreds of different places. Eventually, the file was brute-forced, and, since there was no special salting, the majority of the passwords were revealed. It was easy for Maria to download the file with all passwords in plain text. While skimming it, she saw the password: *john1985service-name*, where *service-name* was the name of the hacked company. Immediately, she recalled that her best friend, John, was born in 1985. Without a second thought, she tried to log in with John's credentials in Twitter. She knew John's username, since she was following him in Twitter, and she thought to try the password: john1985twitter. Voila!

8.3 Key properties

This story highlights a number of security issues which are prevalent today and we expect to have in the near future. We enumerate them here.

- Mobile users are increasing day by day. Many of them do not have access to a desktop PC or laptop. They use all digital services through their smartphone. User interfaces in smartphones are less convenient than ordinary interfaces provided by desktops/laptops. Users are, somehow, coerced to select easy-to-remember, and more importantly, and easy-to-type passwords. These passwords are usally weak.

- Brute-force attacks are considered unrealistic. If you have three login attempts before CAPTCHA solving comes in place then a brute-force attack is highly impractical. However, even major IT companies have security weaknesses, which, if leveraged, can result in password leakage. Attacking an off-line list of passwords using brute-force is not unrealistic anymore, since the cloud evolution and GPU acceleration in cryptographic operations have impressively increased the efficiency of dictionary attacks.

8.4 Final remarks

The goal of Maria's attack was to jokingly convince her friend John to raise his concerns about security. In the same fashion, a determined attacker could hijack many of John's accounts, steal his personal information, and, eventually compromise his digital profile by accessing his e-mail, his e-banking, etc. Users often think of all available services as totally independent to each other. They select hard passwords for sensitive services, such as e-banking, but weak passwords for entertainment sites. Moreover, they trust the security of the sites they visit and they use easy-to-guess patterns for passwords. Users rarely consider the possibility of a large web site being compromised by an attacker, and all users' information, such as passwords, being leaked in the wild. Security, many times, involves many different pieces or layers connected together. The attacker needs only to find the weakest piece, or the layer which can be easily compromised, for escalating his way to the real target. Even large web sites can have vulnerabilities that if exploited can leak sensitive information about their subscribers. An attacker should be unable to get any hints about a user's credentials for a service A, based on leaked information of the user's credentials of service B.

9 Research Roadmap

Contents

As explained in Deliverable *D4.1: First Report on Threats on the Future Internet and Research Roadmap*, the process we adopted to define the initial research roadmap was based on a number of brainstorming activities conducted by the members of the SysSec consortium and several international experts. In particular, to bootstrap the process, we started from the list of future threats identified at the end of the Forward project, and published in the Forward White Book.

Even though the roadmap was specifically designed to cover only a short- to mid-term timeframe, it is very unlikely that large changes to the security landscape would force us to completely shift our directions on a yearly basis. This is, in fact, a desirable characteristic for a roadmap, since it is unrealistic to believe that the research community can completely change its objectives and priorities every few months.

However, the fact that a certain amount of stability is an important feature does not mean that we have to stop refining and improving our roadmap to take into account changes in the system security landscape. In particular, we can identify four main reasons that can lead to modification of the roadmap's directions:

- New threats and attacks are discovered that need to be addressed by the research community.
- Existing threats are mitigated by deployed products, changes in the underlying technology, or new defense mechanisms.
- Existing threats, even if unsolved and still potentially harmful, lose interest because of changes in the underground ecosystem or in the criminal motivations.
- Changes in the existing technology or in the available services suddenly increase the likelihood and severity of some previously unlikely attacks.

In order to make our approach more systematic, we adopted a simple yet effective update procedure. First of all, during the year we keep monitoring several sources looking for new trends in security. This includes scientific papers, statistics about current and future threats reported by antivirus and security companies in their public reports, and opinions of international experts discussed in blogs, talks, whitepapers, or public panels. This helps us to promptly identify new directions that we want to consider as possible candidates for the roadmap. We will discuss some of our findings in the following sections.

In parallel, we also collect information about the research that has been published in the top venues in system security, in the areas identified by our

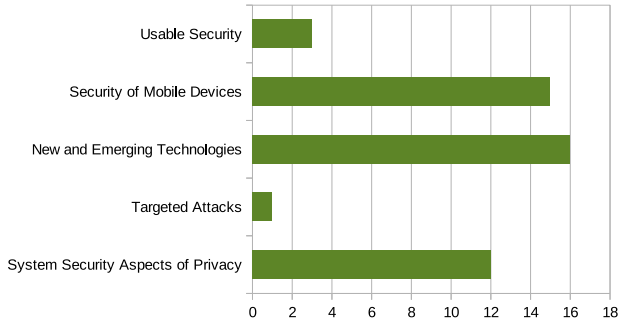

Figure 9.1: Papers published in top tier system security conferences

current roadmap. This is very important because it can provide an immediate feedback on which topics have been covered, which new solutions have been proposed, and which directions still need to be explored.

Finally, we involve a number of external experts in the discussion. This year, this was done in two steps. First, during our working group meeting, we asked each expert to position each threat from the previous roadmap on a number of two-dimensional graphs covering impact, likelihood, R&D priority, etc. This experiment, inspired by the approach adopted to redact the Global Risk 2012 document published by the World Economic Forum, allowed us to support the collected data and to put on a 5-point Likert-like scale [4] the existing roadmap research directions.

In addition, we also organized a panel at the DIMVA security conference regarding the future of malware and underground economy. DIMVA is a premier forum for discussing the advancement in the state of the art in intrusion detection, malware analysis, and vulnerability assessment. During the 2012 conference we ran a panel in which we invited all the attendees to answer a number of questions regarding future trends in system security.

Finally, to conclude our approach, we merged all the collected information and opinions and we summarized and presented the results in a revised edition of the research roadmap.

9.1 Evaluation of the Previous Roadmap

The first step to evaluate the content of our roadmap consisted in reviewing all the papers published in the top system security conferences, to count how many of them matched the directions proposed in our roadmap. We decided to limit our analysis to top-tier venues, to avoid taking into account smaller conferences with potentially narrower focus (e.g., on malware or on

cloud computing) that would have skewed the results toward certain topics. The list of the conferences we took into account is the following:

- IEEE Symposium on Security & Privacy - 2012
- Network & Distributed System Security Symposium (NDSS) - 2012
- ACM Conference on Computer and Communications Security (CCS) - 2011
- USENIX Security Symposium - 2012

Figure 9.1 shows the results of this experiment. Not surprisingly, the security of mobile devices and the one of emerging technologies dominated the publications for the year. On the other end of the scale, usable security and targeted attacks - still considered very relevant topics in the area - were covered by very few scientific papers.

Second, to better formalize and measure the characteristics of the topics proposed in the roadmap, we identified the following six different features:

- *Impact* - represents the severity of the threat as a measure of the amount of damage that can be caused to the users.
- *Likelihood* - measures how likely the threat is to become an issue in the next five years.
- *Research* - is an estimation of the amount of research already done on the field and of the fact that the topic should (or should not) be a priority for the research community.
- *Technological difficulty* - measures how difficult the problem is to solve from a technological point of view. However, it is important to note that *easy* on the technological scale means that a solution would not require a large engineering effort, and not that it would also be easy to deploy, or to convince the users to adopt it.
- *Time scale* - an estimation of how far in the future (up to five years) the threat is likely to materialize.
- *Target size* - represents a measure of how large the percentage of users is affected by the threat.

To simplify the data collection, the six features were grouped together on three graphs: Impact vs Likelihood, Research vs Technological challenge, and Time vs Target scale. Each axis is divided in five intervals, generating a total of 25 square regions on the graph. During the second *SysSec* Working Group meeting, we distributed a number of questionnaires asking external experts in various fields of system security to express their opinion regarding

the six previously mentioned features. We then grouped their preferences and we summarized the aggregated results using a heatmap with the following color code:

BLUE	No votes
GREEN	Few (less than 3) votes
YELLOW	Medium (less than 5 votes)
RED	high (5 or more votes)

In the rest of this section we present the results of our analysis and we report some of the comments extracted from the different experts opinions.

9.1.1 System Security Aspects of Privacy

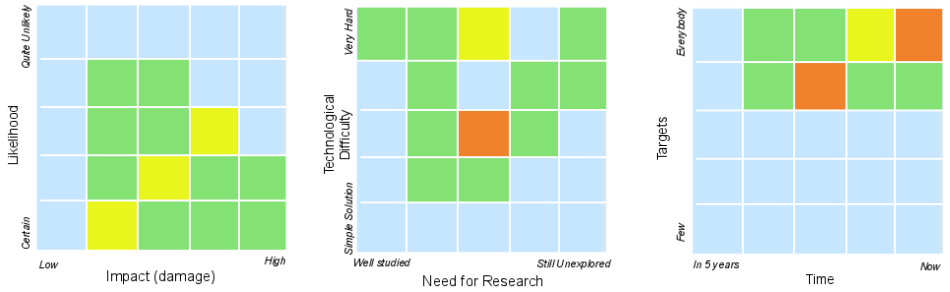

This research direction is focusing on the increasing amount of personal information that is stored online, and in ways in which attackers can exploit them to violate the privacy of millions of users.

The likelihood-impact graph shows a consensus about the fact that this threat is very likely to impact the society, while the experts were dubious regarding the actual impact in terms of damage to the users. The threat is perceived as very hard to solve from a technical point of view, even though the experts recognize that the research community has already done some job in this direction. Finally, everybody agreed on the threat scale: all users will be affected by the problem in the near future.

As we already discussed in Chapter 4.2, during the WG meeting the privacy aspects were largely discussed in connection to the "big data" security problem. In other words, the amount of data available online has increased so much that at the moment we are not even sure we know the right metrics to quantify the threats and risks from the accumulation of so much data. Even worse, we probably need to design new techniques to analyze the collected information to extract interesting relationships between the various data sources.

9.1.2 Collection, Detection, and Prevention of Targeted Attacks

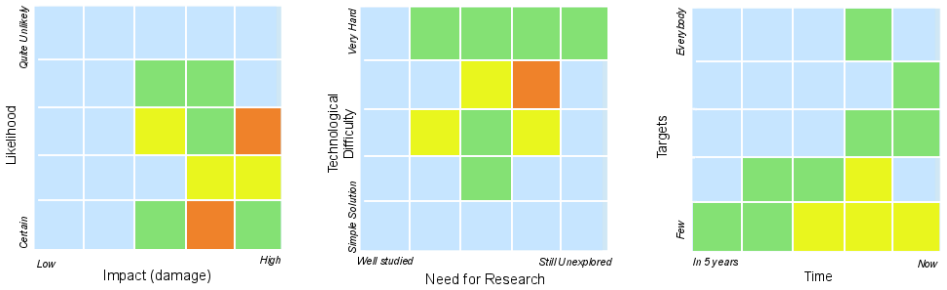

The first graph shows a very clear picture, with most of the experts opinions converging to the 4th quadrant (lower-right) that identifies certain threats that can cause a large damage to society. A clear consensus was also reached for the technical challenge (quite hard to solve) and for the target scale (limited to few individuals). However, the exact time frame and the need to perform research in this field are not clear.

One comment that was often raised during the discussion is the fact that *targeted attacks* still lack a precise definition, and almost everybody has his own personal way to look at the problem. This may be one of the reasons behind the lack of consensus on the research direction.

Another point that was discussed during the WG meeting is the relationship between targeted attacks and social engineering. The common point of view that there is no technological solution to stop social engineering attacks may be incorrect. Certainly, more research is needed to study and mitigate the social engineering threat, that is so often used as the basis for many different kind of attacks.

9.1.3 Security of New and Emerging Technologies

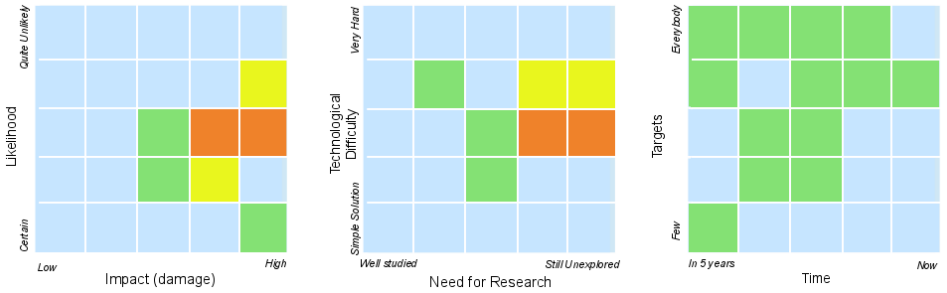

From the experts' votes we can distill two clear messages: 1) the security of emerging technologies will have a large impact, and 2) at the moment it has not been sufficiently explored by researchers. However, our analysis of the yearly publications shows a different picture, with the security of emerging technologies being one of the most represented topic in scientific publications. This discrepancy may be the consequence of the fact that "emerging technologies" is probably too vague a term. In its original description, the term included social networks and smart meters, as well as cloud computing and SCADA networks.

This large spectrum of topics were obviously perceived in different ways from the experts. In fact, most of them chose a mid-range value for likelihood and technological challenge, and a seemingly random distribution for the target-vs-time graph (e.g., social networks are already a problem now, while SCADA networks are more likely to affect us in the near future).

9.1.4 Security of Mobile Devices

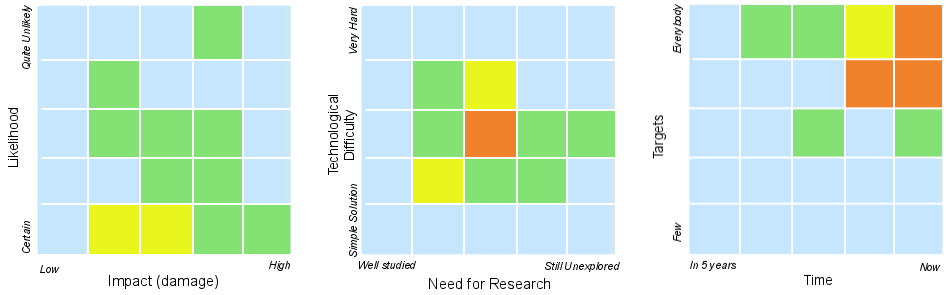

The likely-vs-impact graph is quite scattered, showing the expert perplexity regarding the likelihood and the possible damage to society of mobile threats. At the same time, all voters agreed on the fact that this threat is already happening now and that it affects everybody. A possible way to interpret this contradictory result is, for example, the fact that everybody agrees on the existence of mobile malware but both its sophistication and its damage to the users are still limited, and it is still not clear how such malware will evolve in the near future.

The research-technology graph shows a nearly perfect distribution around the center. The general opinion during the WG meeting was in fact that a large amount of papers have already been published on the topic, but a solution is still missing, since most of the time the user turns out to be the weakest point in this kind of threats.

The working group on malware and fraud also proposed a new threat related to mobile devices: portable malware. This term describes the fact

that recent OS platforms provide an additional advantage for malware authors. In fact, the same kernel is used among different target devices. iOS, for instance, is designed for both iPhone and iPad alike. Some even work on MacOS. Therefore, also malware written for these platforms reaches a broader mass. The same is true for Android, which runs on mobile phones, tablets and even laptops (see Chapter 2 for more information).

9.1.5 Usable Security

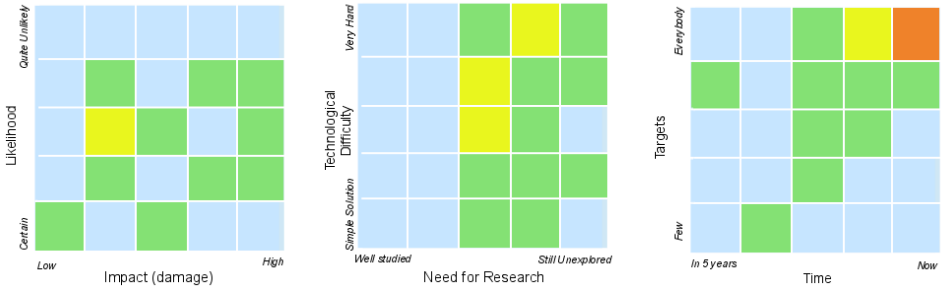

The importance of studying the usability of security solutions is a well-understood problem. However, when experts were asked to put their preferences on a scale, we obtained a large number of different opinions. Despite this general lack of consensus, there was an agreement on two important points: the need to do research in this field, and the fact that this is an urgent problem that affects the entire society.

9.1.6 DIMVA Panel

During the 9th Conference on Detection of Intrusions and Malware & Vulnerability Assessment (DIMVA 2012) we organized a panel which focused on two important questions related to our roadmap: *"Will malware still be a problem in 5 to 10 years from now?"* and *"Will cybercrime remain economically feasible?"*.

The participants were given colorful post-it notes and asked to write their opinion on them. Then, the notes were posted on poster boards and the PC chair led the discussion commenting and inviting comments on the post-it notes. From the answers we collected from the attendees, we can distill two important points. First, malware will certainly be a problem in the future, but probably it will evolve to a different form. Low-effort, mass-market malware will decrease in prevalence and more sophisticated malware will increase due to the involvement of government actors. At the same

time, malicious code will move from traditional computer systems toward the smart-home/device ecosystem (videogame console, tv sets, cars, ...). Most of the experts agreed on the fact that the current reactive approach adopted by security companies is not likely to solve the malware problem in the near future; things will get worse before they get better.

The second point we can take away from the panel discussion is that cybercrime will remain profitable in the future unless we come up with radically different approaches to deal with the trade-off between security and usability. At the same time cybercrime will become more professional and more players will emerge in the field. To mitigate the problem we have to limit the low hanging fruit (i.e., the simple and less risky attacks), try to increase the cost of performing malicious activities, and push to have modern cybercrime laws in as many countries as possible.

9.2 Updated Roadmap

The results of our internal study and the opinions collected from international experts that are not part of the SysSec consortium confirmed the importance and the significance of our roadmap. The directions we proposed at the end of the first year of the project are still relevant today, and no major change has been proposed for this revised version. However, some of the topics have been either refined or focused to more specific problems, resulting in different research recommendations.

9.2.1 Big Data Security & Privacy

More and more personal information about an increasing number of users will be stored online in the near future. Social networking sites are a very well-known example of this trend, but, unfortunately, they are just the tip of the iceberg of a much larger phenomenon. File hosting services, cloud computing, back-up solutions, medical databases, and web emails are other examples of services that store personal information outside the direct control of the users.

Such a large amount of information requires to be carefully protected and regulated in order to preserve the citizens' privacy. This includes several complementary aspects: cryptography to store the data, system security to properly protect the data from being stolen, data mining and correlation to understand the hidden connections between data sources, and new mechanisms to detect impersonation and identity thefts.

Recommendations and Research Directions:

> Researchers should investigate how to protect users against sophisticated attacks that aim at disclosing their personal information or stealing their identities. New research is also needed to develop automated and scalable techniques that can be applied to big data sources.

Expected Impact

- Increased confidence by EU citizens in a privacy-preserving use of ICT.
- Increased societal acceptance of ICT through the assured protection of basic privacy expectations.
- Increased support towards the protection of the right of privacy for ordinary citizens.

9.2.2 Targeted Attacks - The Needle in a Haystack

Targeted attacks are still a priority in the second version of our research roadmap. These attacks clearly showed how our current defense tools, policies, and infrastructures failed in front of a threat that was designed to focus against a specific target instead of blindly targeting the entire community.

Last year we mentioned malicious hardware as an example of sophisticated attacks that can be used to target high-profile organizations, and for which there is still little or no defense.

This year we want to mention a more traditional, less technical, form of attack: social engineering. Social engineering consists of manipulating people to get access to private information or computer systems. It is often used in combination with other form of exploitation to perform targeted attacks. Moreover, since social engineering targets humans instead of the computer systems, it is very hard to deal with from a technological point of view.

Recommendations and Research Directions:

> We believe it is very important for researcher to develop new techniques to collect and analyze data associated to targeted attacks. The lack of available datasets, in addition to the limitation of the traditional analysis and protection techniques, is one of the current weak points of the war against malware. In this area, the problem is often to find the needle of the targeted attack in the haystack of the traditional attacks perpetuated every day on the Internet.

A second important aspect that needs to be investigated is the relationship of targeted attacks and social engineering. In particular, new defense techniques need to be proposed to mitigate social engineering attacks.

Finally, researchers should also focus on new defense approaches that takes into account alternative factors (such as monetization), and large scale prevention and mitigation (e.g., at the Internet Service Providers (ISP) level).

Expected Impact

- Significant improvement towards the protection of Critical Infrastructures.
- Winning significant ground against sophisticated cyber attackers.
- Design of new detection and protection techniques to mitigate cyberespionage attacks against governments and large organizations.
- Improved collaboration with international research and operational stakeholders.

9.2.3 Security of New and Emerging Technologies

Analyzing and securing emerging technologies has always been a core objective in the area of system security. Unfortunately, it is often the case that new services and new devices are released before the research community had a chance of studying their security implications.

Last year we proposed four topics, in the area of new and emerging technologies, that needed to be urgently studied from a security point of view: Cloud Computing, Social Networks, SCADA Networks, and Smart Meters. Even though these four topics are still relevant today, this year we want to focus on a particular aspect of the problem: the lack of tools and techniques to investigate most of these technologies. In fact, as we explained in Chapter 3.3, due to their black box implementation and proprietary nature, several environments are hard to analyze by researchers. For instance, this is true for the smart grid, for SCADA networks, for automotive systems, and to some extent also for more common systems like cloud infrastructures.

The availability of open source tools, specifications, and techniques so that researchers can properly analyze the emerging technologies is a key factor for the success of system security in this area.

Recommendations and Research Directions:

Security of new and emerging technologies before it is too late is one of the main priorities of the system security area. In this direction, it is important to sponsor activities and collaboration between academia and the industrial vendors to maximize the impact of the research and reduce the time required for the analysis and the experiments.

Vendors should also help in the development of the required tools to allow researchers to analyze the system and study their security.

Expected Impact

- Increased adoption of, and placing trust in, emerging technologies by ordinary citizens.
- Reduced costs associated with security incidents.
- Lower barriers for mobile operators and application developers to provide accessible and affordable mobile services to their customers.

9.2.4 Mobile and Smartphone Security

We are currently witnessing the penetration of mobile devices in every facet of our society. Exploiting such devices is often easy due to a number of factors such as the limited computational power to run full fledged security software, the dependency on battery power that may make it unpractical to run a software for a long period of time, the lack of security design, and the focus on usability at the expenses of security.

However, large-scale epidemics of malware and attacks against mobile devices are yet to be encountered. This could also be a consequence of the fact that, currently, all mobile devices use a market to provide applications for the user. Since these markets are supervised and applications are rated by users, it is harder for an attacker to distribute malware over these channels (see Chapter 2.2 for more details).

However, things can change quickly. For example, the upcoming release of Windows 8 and its mobile version can open the door to more portable malware samples to reach a broader community.

Recommendations and Research Directions:

In the past year researchers have published several papers on the topic. However, we still believe that we need more research focused on the development of defensive tools and techniques

that can be deployed to the current smartphone systems to detect and prevent attacks against the device and its applications.

Expected Impact

- Increased adoption of mobile devices for commercial use by ordinary citizens.
- Improved European industrial competitiveness in mobile phone applications in all realms of life.

9.2.5 Usability

Research on usable security is now mostly relegated to few experts and a couple of dedicated conferences. Some researchers believe that security can be made user friendly and users can be instructed to properly behave to reduce the risks. Others think that the only way to secure a system is to remove any potentially dangerous choice, moving the decisions from the users to the system itself. Either way, the usability of security will play an important role in the next few years.

The impact of new defense techniques greatly depends on the assumption made on the final users and on their involvement in the security process. Unfortunately, it is also a very difficult problem to solve.

Recommendations and Research Directions:

We believe that a study of the usability of security countermeasures is very important and it will become even more critical in the future. If we want to progress in this direction, we need *interdisciplinary* efforts that bring together experts from different fields (engineering, system security, psychology, ...).

Expected Impact

- Empowering users to play a more effective role in securing cyber space.
- Provide increased support to end users so as to make better decisions when accessing the ICT infrastructure.
- Increase the end-user adoption of security-related software and monitoring systems.

Bibliography

[1] Theodosis Dimitrakos, Fabio Martinelli, Peter Y. A. Ryan, and Steve A. Schneider, editors. *Formal Aspects in Security and Trust, Fourth International Workshop, FAST 2006, Hamilton, Ontario, Canada, August 26-27, 2006, Revised Selected Papers*, volume 4691 of *Lecture Notes in Computer Science*. Springer, 2007.

[2] A.P. Felt, M. Finifter, E. Chin, S. Hanna, and D. Wagner. A survey of mobile malware in the wild. In *Proceedings of the 1st ACM workshop on Security and privacy in smartphones and mobile devices*, pages 3–14. ACM, 2011.

[3] D.G. Feng, M. Zhang, Y. Zhang, and Z. Xu. Study on cloud computing security. *Journal of Software*, 22(1):71–83, 2011.

[4] Eleni Gessiou, Alexandros Lambrinidis, and Sotiris Ioannidis. A Greek (Privacy) Tragedy: The Introduction of Social Security Numbers in Greece. In *8th ACM CCS Workshop on Privacy in the Electronic Society (WPES)*, 2009.

[5] Eleni Gessiou, Stamatis Volanis, Elias Athanasopoulos, Evangelos Markatos, and Sotiris Ioannidis. Digging up social structures from documents on the web. In *IEEE 2012 Communication and Information System Security Symposium (Globecom CISS)*, 2012.

[6] P. Gilbert, B.G. Chun, L.P. Cox, and J. Jung. Vision: automated security validation of mobile apps at app markets. In *Proceedings of the second international workshop on Mobile cloud computing and services*, pages 21–26. ACM, 2011.

[7] Daniel Halperin, Thomas S. Heydt-Benjamin, Kevin Fu, Tadayoshi Kohno, and William H. Maisel. Security and privacy for implantable medical devices. *IEEE Pervasive Computing*, 7:30–39, 2008.

[8] M. Jensen, J. Schwenk, N. Gruschka, and L.L. Iacono. On technical security issues in cloud computing. In *Cloud Computing, 2009. CLOUD'09. IEEE International Conference on*, pages 109–116. Ieee, 2009.

[9] Alexandros Kapravelos, Iasonas Polakis, Elias Athanasopoulos, Sotiris Ioannidis, and Evangelos P. Markatos. D(e—i)aling with VoIP: Robust Prevention of DIAL Attacks. In *ESORICS*, pages 663–678, 2010.

[10] Samuel T. King, Joseph Tucek, Anthony Cozzie, Chris Grier, Weihang Jiang, and Yuanyuan Zhou. Designing and implementing malicious hardware. In *Proceedings of the 1st Usenix Workshop on Large-Scale Exploits and Emergent Threats*, pages 5:1–5:8, Berkeley, CA, USA, 2008. USENIX Association.

[11] Martin Mulazzani, Sebastian Schrittwieser, Manuel Leithner, Markus Huber, and Edgar R. Weippl. Dark clouds on the horizon: Using cloud storage as attack vector and online slack space. In *USENIX Security*, 8 2011.

[12] ZDNET report. Hackers steal keyless bmw in under 3 minutes. http://www.zdnet.com/hackers-steal-keyless-bmw-in-under-3-minutes-video-7000000507/, July 2012.

[13] S. Skorobogatov and C. Woods. Breakthrough silicon scanning discovers backdoor in military chip (draft of 05 march 2012).

[14] Salvatore Stolfo, Steven M. Bellovin, Angelos D. Keromytis, Sara Sinclair, Sean Smith, and Shlomo Hershkop, editors. *Insider Attack and Cyber Security: Beyond the Hacker (Advances in Information Security)*. Springer, 2008.

[15] W. Zhou, Y. Zhou, X. Jiang, and P. Ning. Droidmoss: Detecting repackaged smartphone applications in third-party android marketplaces. In *Proceedings of the 2nd ACM Conference on Data and Application Security and Privacy*.