SEVENTH FRAMEWORK PROGRAMME

Information & Communication Technologies
Secure, dependable and trusted Infrastructures

COORDINATION ACTION

Managing Emerging Threats in ICT Infrastructures

Grant Agreement no. 216331

Deliverable D2.2:
Second workshop report

Contractual Date of Delivery	31/05/2009
Actual Date of Delivery	17/06/2009
Deliverable Security Class	Public
Editor	Engin Kirda
Contributors	FORWARD Consortium
Quality Control	Christopher Kruegel

The FORWARD Consortium consists of:

Technical University of Vienna	Coordinator	Austria
Institut Eurécom	Principal Contractor	France
Vrije Universiteit Amsterdam	Principal Contractor	The Netherlands
ICS/FORTH	Principal Contractor	Greece
IPP/BAS	Principal Contractor	Bulgaria
Chalmers University	Principal Contractor	Sweden

Contents

1 Introduction **5**

2 Working Group Discussion Summaries **7**
2.1 Session: Malware and Fraud 8
2.2 Session: Critical Systems 10
2.3 Session: Smart Environments 13

3 Position Papers **19**
3.1 Workshop Plenary Talks 19
3.2 Workshop Work-In-Progress Talks 28

4 Conclusions **33**

5 List of Participants **35**

CONTENTS

Chapter 1

Introduction

This deliverable summarizes the activity of the second FORWARD workshop. This workshop constituted the end of the second phase of the project. The aim of this second phase was to establish a number of working groups; each working group had to identify a number of emerging threats in their respective areas (malware and fraud, smart environments, and critical systems). These threats were summarized in three threat reports (Deliverable D2.1.x), one per working group. The goal of the second workshop was to checkpoint and critically review the work that has been done in the working groups, in particular, the threat reports. More precisely, each working group should present their threats to a larger audience comprised of experts. In discussions and presentations, we wanted to make sure that the lists of threats are comprehensive – that is, each working group has identified all major threats in their respective areas. Moreover, we wanted to use the workshop to establish an initial ranking for the threats presented by each working group. Clearly, at one point, it is necessary to prioritize threats and focus the attention on those that present the largest threat potential to ICT infrastructures and the society at large. Of course, the assessment of the danger that each threat poses, as well as an analysis of inter-dependencies among threats, is a focus of the third project phase (which is to be completed by the end of the year). However, we attempted to leverage the presence of a large amount of domain experts to obtain an initial ranking that would combine and reflect the viewpoints of a large audience.

For the second workshop, we decided to invite a number of selected speakers that would give presentations at the beginning of the workshop on the first day and later during the second day. The talks set a framework in which the detailed technical discussions about the individual threat reports could take place. For these discussions, the attendees would first break into working group sessions to perform the necessary review of the threats that each group had defined. Then, in a next step, the outcome of each discussion was presented to the audience at large. This two-step process served two purposes. First, in the actual discussion sessions, we had less people involved. This made the discussion process manageable and interactive. In the second step, we presented our findings in a succinct fashion to the whole

audience. This allowed everybody who participated in the first discussion round to ensure that their opinions were correctly reflected. In addition, it allowed people that were present in other working group discussions to see what other groups did, and to provide feedback.

According to Annex 1, a total of 60 attendees was considered to be the threshold for a successful workshop. This threshold was significantly exceeded, with a total of 103 attendees. This clearly demonstrates the significant interest and participation to the FORWARD working groups and workshops. Moreover, non-academic participation remains to be strong. 39 attendees (37.8% of the participants) came from industry or policy-making institutions.

In this document, we first summarize the three working group discussions that were held during the two-day workshop. In addition to the discussion sessions, a total of 11 talks were given in the form of plenary talks and keynotes. Moreover, we had 7 five-minute work-in-progress talks. These talks are summarized in the subsequent chapter. Finally, we discuss the conclusions that the consortium has drawn from the workshop, and we briefly outline the future actions that we plan to take in the subsequent, third phase of the project.

Chapter 2

Working Group Discussion Summaries

In the afternoon of the first day of the FORWARD workshop, the participants split up into three parallel tracks – one for each working group. The split was done based on each participant's interests and expertise. The three resulting groups had parallel sessions, and joined a discussion on the future and emerging threats for each domain handled by each working group. The target of the discussion was: (a) receive feedback on the future threats that the working groups have already identified, (b) extend the list of possible future threats with the ones envisioned by the participants, and (c) classify threats based on their importance as perceived by experts of the community.

As a quick reminder, the topics and focus of the three working groups are as follows:

The *Malware and Fraud* working group is concerned with the malware and fraud-related threats on the Internet. It covers topics that range from novel malware developments over botnets to cyber crime and Internet fraud.

The *Critical Systems* working group focuses on critical systems whose disruption of operation can lead to significant material loss or threaten human life. It attempts to identify emerging threats in this area.

The *Smart Environments* working group is concerned with ordinary environments that have been enhanced by interconnected computer equipment. There is general expectation that a large number of small devices such as sensors and mobile phones will be interconnected. The group aims to identify emerging trends with respect to security in this domain.

In the following three sections, the findings and conclusions of each working group discussion meeting are summarized.

2.1 Session: Malware and Fraud

During the afternoon session of the second FORWARD workshop, the working group for *fraud and malware* was confronted with a controversial issue. Guided and coordinated by Engin Kirda and Christopher Kruegel, the main objectives were to review previously identified threats and to create a relevance rating for them. The initial plan, which was to create a ranking from 1 to 10 (depending on the threat's perceived relevance for the future) failed. The reason was that even the experts in particular fields like malware authoring or social networks could not appoint such a concrete rating to one of the topics. Even with a lot of knowledge and information, a prediction of how relevant a specific topic will be in 10 years from now is hard to give.

Instead, the participants decided to take a different approach, and categorize the identified threats in three different categories, depending on their importance. To assess these *threat level*, different metrics were identified and applied during the discussion. The three most relevant metrics are:

1. **probability of occurrence:** This metric describes the participant's assessment of the probability, that the attack in question is actually carried out. Attacks with malicious hardware for example, are very likely to yield a positive result for the attacker, if carried out properly. However, the required means to carry out such an attack in the first place, reduces the probability of such an attack enormously. Other factors for the probability are motivation of the attacker and possible gain in case of a successful attack.

2. **impact:** Not every attack influences the whole Internet community when unleashed. Therefore, an impact rating for specific threats is obligatory. The impact describes, how many users are affected and what damage level is to be expected. A worm like *Conficker* for example, affects millions of users at the same time, while the damage it causes is less serious than other worms seen on the Internet.

3. **relevance:** The last measurement that was produced during the discussion concerns the relevance to the Internet community as a whole. Not every fraudulent or malicious action necessarily belongs to the IT domain. Therefore, the relevance for this domain also influences the perceived risk connected to a specific threat.

With this weighting mechanism, the working group was able to conduct a productive discussion. With all threats that were identified previously, Figure 2.1 visualizes the outcome of the discussion and the assigned ranking for each element.

Some threats were thoroughly discussed during the session, because expert opinions were rather controversial. The following sections discuss these specific points.

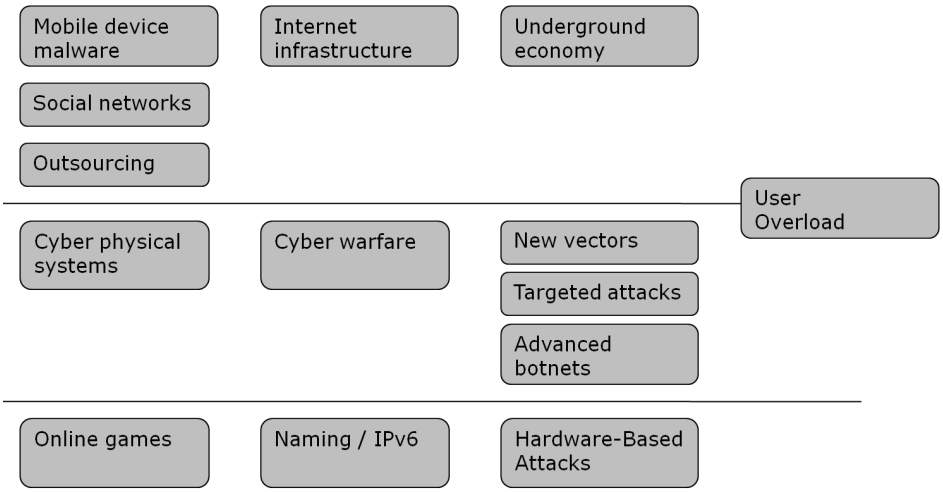

Figure 2.1: Threat ranking.

2.1.1 Social Networks

Social networks like Facebook recently experienced an enormous boom. From a technical perspective, they are not more dangerous than every other web site. The difference simply lies in the amount of private and. possibly, sensitive data that is represented by a user account. Research has shown, that most of the participating users are less suspicious within the community than, for example, to spam-mails. If this heightened level of trust diminishes to the same level as for the rest of the Internet remains to be seen.

2.1.2 Underground economy

Although strictly spoken not a direct threat, the underground economy was identified as the ultimate enabler for various sorts of attack. From spam campaigns to credit card fraud, money laundry and rent-a-botnet operations, the underground economy always represents a major part of the basic attack vector. Put together and interconnected, the single technologies and exploits form an economy that is based on the same principles as the common economy: **supply and demand**. As a result, the participants within a spam campaign, for example, can be manifold, forming a group of involved individuals with different objectives and motivations. The working group agreed that this structure will also apply to the future. Therefore, it is rated as an important threat in the final result

2.1.3 User Overload

A lot of today's threats like cross-site scripting, phishing or similar attacks could be mitigated by various techniques. In reality however, these solutions often require a

user to have at least a certain understanding of what the threat means to him. Furthermore, this knowledge is important to let a user decide what a specific warning dialog means. Even today, users become ”resistant” to dialogues. They strongly tend to get rid of annoying interruptions by clicking OK on each appearing question. This problem is not a technical one, but of the user interface. Nevertheless, it is imperative to wrap new solutions to upcoming and even existing threats in understandable and discreet user interfaces to make sure, they are properly used. The working group agreed, that this overload is a constant problem that is very likely to persist for a long time and hinder solutions for security problems to catch, even if they already exist.

2.2 Session: Critical Systems

Erland Jonsson coordinated the session for the Critical Systems Working Group (CS WG). The objectives of the parallel sessions were to collect direct feedback about the emerging threats that had been identified in the first workshop and then further refined through discussions within each working group. The FORWARD consortium especially wanted to discover what a wider group of experts thought about the *relevance* of the identified threats, if the list can be seen as *complete* or whether some significant threats are missing. As input to the discussion of the final result, we also asked about a ranking of the presented threats. More than 35 people chose to attend the Critical Systems Working Group session.

The agenda for the CS WG session was as follows. First, Erland Jonsson was to present the background of the working group, followed by each of the emerging threats that had been identified by the working group. It was decided that each threat should be presented one at a time, and that the audience could give feedback directly. When all threats had been presented, there was a time slot for giving overall feedback on the work and comment on any potentially missing threat. Three people also expressed a willingness to give a short presentation on material related to CS WG. Damiano Bolzoni was to present the work done on SCADA systems. Aljosa Pasic was to present the results of the first workshop of the PARSIFAL project. Finally, Hong-Linh Truong would present his thoughts on how to emphasize the human role in the critical system.

After the agenda was settled, Erland Jonsson started presenting the background of the CS WG. The working group was formed during the first FORWARD workshop in Goteborg, Sweden, and has since then compiled and discussed important threats related to critical systems. Erland Jonsson spent some time describing the focus of the working group and its scope. This can best be described by Figure 2.2. He emphasized the fact that *new emerging critical systems* exist, such as the *connected car*, and that systems such as these also need to be considered in the current work.

A question was asked from the audience related to our use of the words *security* versus *safety* (see Figure 2.2). Several people in the audience chimed in to clarify

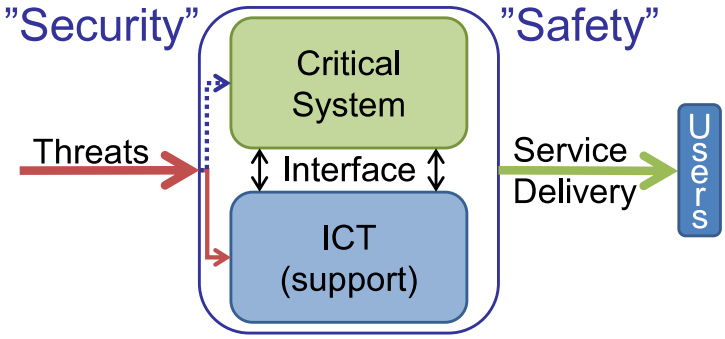

Figure 2.2: A model of a specific system for critical services. We do not consider threats directly targeting the critical system (CS) (dotted line).

the meaning. Overall, people agreed with our defined scope of the group and the use of the terms.

After the introduction, Erland Jonsson presented the list of emerging threats that had been identified by the experts connected to the working group. At this time, the threats were only sorted into the methodological axes the FORWARD consortium has identified for guiding the work.

When presenting *Sensors as the "New Computing Class,"* Erland pointed out that in such environments the adversary may sometimes have more powerful hardware than what exists within the sensor network under attack. A person in the audience commented that we should also consider that the attacker may have more powerful software.

The following threat related to New Generation Networks (NGNs) led to a wide discussion. First, the audience asked about our definition of *threat* and Erland Jonsson had to reiterate some of the earlier discussions we have had within the working group to set a common baseline for all the participants in the session. Fortunately, we also had copies of the draft report produced by the working group, and we could there show the threat definition used by the three working groups. Having settled the background, the discussion turned to the role of the hardware versus the protocols. There was expressed an opinion that security stays in the applications and they should be made more secure. Michael Behringer had a comment regarding the complexity of networks. He believes that systems are becoming so complex that they cannot really be understood and this in turn will lead to more insecurity. Some people wanted a more general threat description regarding the infrastructure running the network, and not towards NGNs specifically. Some participants also voiced the concern of the cost considerations that is now often governing decisions. It is seldom people are willing to pay for resilience, even in the critical systems' domain.

There were minor comments to the following three threats presented, but the next topic that spurred a larger discussion was the *Use of COTS components*. One of the problems here is that different groups of professionals define *COTS* differ-

ently. Aljosa Pasic summarized the viewpoint of several people in the audience, when he pointed out that we should emphasize two points when discussing the risks related to COTS components and systems. The first danger is the use of COTS components in a context never envisioned for them. The second is the compositional effects that have never been tested when many COTS systems are working together.

After a break, threats related to the human factor were presented. Many in the audience agreed with the descriptions of the problems related to the human factor and some even went as far as stating that threats related to this group are among the most important. As systems become more complex, users have a larger ability to make more severe mistakes. There was a consensus that technology that *adapts to the human* is very important but many times forgotten. Humans will make mistakes and the technical system should compensate for them. Furthermore, the human factor has been identified in security research for a long time, but the area is not very well explored. There should be more research efforts focused on this area.

After the threats were presented, Erland Jonsson asked for feedback from the audience to create a first ranking of the threats. Among the threats that were seen as the most important, several corresponded to the human factor. The result was the following informal order among the threats:

- *Threats deemed as very significant*
  - The Human Factor
  - The insider threat to critical infrastructures
- *Threats deemed as significant*
  - Wireless communications in critical industrial applications
  - Retrofitting security to legacy systems
  - Sensors as the “New Computing Class”
  - New Generation Networks
  - Hidden functionality
- *Threats deemed as less significant*
  - Use of COTS components
  - Safety takes priority over security
  - Cultural differences between control and security communities
  - Attacks against the Internet

We realize that this simple inquiry is not scientifically well founded, but could still serve as a first approximation.

After the discussion, Damiano Bolzoni presented his work under the project called ”S3SCADA – Secure and Survivable SCADA”. There were a number of

questions related to his research, which led to quite a lively discussion. After that Aljosa Pasic summarized the first workshop of the PARSIFAL project. This is a European Union project closely related to the Critical System Working Group, which deals with the financial infrastructure protection and it was very interesting to hear about the findings of this workshop. Finally, Hong-Linh Truong spoke about the role of the human in critical systems, and the need to make the human more explicit when we discuss such systems.

After these discussions and presentations, the working group session was closed.

2.3 Session: Smart Environments

Coordinator of the smart-environments working group was Sotiris Ioannidis (SI) representing ICS-FORTH, and he was joined to lead the discussion by Georgios Portokalidis (GP) representing VU Amsterdam. The audience consisted of well known members of the research community from all over the world. To kick off the session, SI gave a short presentation about the FORWARD project and the smart-environments group in particular. The scope of the smart-environments group, and a list of future threats as generated by the project were also presented. The token was then passed to the audience and discussion was initiated.

The second day of the workshop, the leader of each working group gave a presentation discussing the conclusions made from the previous day's sessions. This arrangement made possible for workshop participants that were allocated in one of the other parallel sessions to also provide their opinion on the subject.

In the rest of this section, we will discuss the feedback and insights gained during the workshop . We will close with a rating of each threat as it was expressed by the participants of the initial group, and any comments received afterwards.

2.3.1 Session Feedback

Smartphones. Smartphones have been identified by FORWARD as one of the major vector of new threats in smart-environments. The argument is based on the facts that smartphones today are very similar to PCs in many aspects, but they also exhibit special traits as: high-connectivity on different networks (3G, WiFi, Bluetooth and others), limited power (battery), lack of security software.

During the discussion the audience gave special notice to all the sensors that are being used on smartphone devices. Examples of sensors that were mentioned are: GPS, compass, accelerometer, camera, and microphone. These sensors generate valuable information for the user of the phone, such as location, movement direction, video and audio. Furthermore, these sensors are currently controlled fully by software, which means that they can be easily exploited by attackers having compromised the device. The user of the phone receives little feedback on the sensors that are activated at any given moment. A counter example from the PC world was brought up. Cameras embedded on notebooks these days included a hard-wired

notification LED, indicating whether the camera is activated. The problem is exacerbated by the vendors of these devices, who are actually moving in a different direction, trying to hide as much complexity of their system as possible.

Sensors. Besides the sensors we just mentioned above, today we are surrounded by a great number of sensors designed to make our environment "smarter". Such sensors might include for example temperature, and movement sensors located in many modern offices.

Some of the participants thought that it is a problem that is frequently ignored due to the simplicity of the sensors, and the information generated. An example showing that even these sensors can become a threat was brought up, mentioning how researchers were able to use temperature sensors to detect when a certain person enters his office. They were able to deduce this information by looking on small temperature increases that occur when a person is within a room.

RFID. RFID chips are today continuously increasing in usage. Researchers have already shown that there are ways to attack such chips by: spoofing, skimming or relaying.

A possible important issue brought up during the session was the possibility of impersonation of a legitimate user by an attacker.

Home and Office Automation. An increasing number of devices in our home and office is today interconnected, in an attempt to automate many daily mundane tasks for users. For example, consumer electronics such as fridges, photo frames, clocks, etc, are given wireless networking capabilities to enable to talk to each other or a central server.

An important issue identified during the session was that users are frequently unaware of this. This introduces potentially malicious vectors within a home network that users are completely unaware of. As is usually the case, unexpected attacks can often cause cause more damage. Another point also made, was that offices are far more likely to be attacked, as there are far larger potential financial gains to be made.

Social Networking. Social networks have been growing at an amazing pace. Users tend to share private information on these networks without considering the potential impact this data may have in the wrong hands.

During the session a participant mentioned how such networks are used by employers in the US. By going through an applicants Facebook profile, employers attempt to extract information that they use to help them decide on hiring. As users are often careless, or do not even control the data that enter these networks, important privacy issues arise.

Aviation Security. Aviation security is obviously a very important subject, and in many ways overlaps with the work done by the critical system working group. In our case we examine the threats that can be introduced by making the environments within an aircraft smarter. Modern aircrafts are fitted with multimedia systems, and offer services such as telephony and Internet access.

Some very serious concerns were brought up by members of the session. Firstly, it was mentioned that because aircraft makers attempt to reduce weight as much as possible, the fibre installed in the aircraft to deliver control messages, and sensor information is actually shared with the entertainment system. Furthermore, makers are already considering using wireless network to connect computers within their aircrafts to further reduce weight. These facts may introduce important security problems. Attackers could be able to exploit faults in the entertainment system to tap into the aircraft's control, or falsify sensor data. Moving into wireless communications, poses an even greater threat as it would make it potentially easier for attackers to tap on the network, or could enable them to jam communications.

Vehicular Automation. Vehicle automation is another area that the smart environments group overlap with the critical systems group. During our session we discussed issues that arise from vehicle communication with its environment.

A participant mentioned that such systems are already used. For example, in Valencia, Spain ambulances are fitted with sensors that communicate with the traffic lights system to ensure a fast path to the hospital without red lights. It is obvious that the abusing systems such as this consists an important threat.

Multicore Threats. Multicore systems are becoming the standard today. This parallelism might give ground to new threats that aim to exploit race conditions and interdependencies in this new technology.

Discussion on this issue concluded that it is probably very hard for attackers to exploit such issues. Even though, such attacks would affect a significant size of the population, trends that show that attackers pick the path of least resistance to accomplish their goals, make it improbable that such issues will be extremely important.

Malicious Hardware. As most hardware fabrication is nowadays outsourced, malicious hardware is also becoming an issue. Circuits can be added on chips at the fabrication plant to offer a backdoor to potential attackers, or perform some other action. It is technically very hard for vendors to detect whether the produced hardware follows their design to the letter.

Members of the session agreed that this is becoming an important issue. Specially, in the US the possibility of malicious hardware used for espionage, or even for terrorist activities is considered an emerging threat. Potential solutions to this problem were discussed, with the most prominent being the use of secure and trusted fabs for critical hardware such as the one used in aviation and the military, and the runtime detection of such malicious hardware.

2.3.2 Threat Rating

During the session a rating of the various threats was starting to emerge. To assist the procedure of assigning a criticality level to each threat, SI proposed that the session reach a consensus on the factors that play the most important role on determining this rating.

The following factors were identified:

1. *Immediacy:* This factor determines how imminent a threat is.
2. *Awareness:* This factor determines the level of awareness of the public and the research community for each threat. The higher the awareness the more likely it is that the threats will be addressed in a timely fashion before they are realised.
3. *Impact:* This factor determines the size of the affected population by a certain threat.

Having defined a set of factors to assist the classification of threats, the group proceeded to classify threats in three ranks: *High*, *Medium*, and *Low*.

High Ranked Threats. It was agreed that threats that immediately impact the privacy of users and organisations are the most critical ones. Such threats can have a significant and immediate impact to a large number of users. Privacy related threats dominate the areas of:

- Smartphones
- Sensors
- Home and office automation
- Social networks

The problem can be only made worse, if one could aggregate the information from all this areas. Public awareness seems to be a very important factor on ranking privacy related threats the highest. Today, most users seem to be either unaware, or to be ignoring the gravity of such threats. It is necessary then to look at these issues seriously and promptly.

Medium Ranked Threats. Threats that we will encounter frequently in the future were included here. Such are:

- Vehicular automation
- Aviation security

Threats in these areas are very real, and it is very likely that we encounter them in the future. The fact that the industries involved in these areas can comprehend the seriousness of potential threats allows us to classify them as of *medium* importance.

We should note here that during the discussion on the second day of the workshop, some objections were expressed on whether the threats belonging to this ranking should not be more important. The argument on this being that such threats seem to have larger impact. Even though we do partially agree, we do also do believe that industries in this area are more mature when addressing the issues.

Low Ranked Threats. Threats that are very hard to realise, or sometimes their solution relies to policy making were put in this category. Such are:

- Multicore threats
- Malicious hardware

Threats in these areas should not be considered trivial to defeat, or unimportant. Nevertheless, having a finite amount of resources to use for tackling potential future threats, our discussion with some of the research community's experts showed that these should be of lower priority.

Chapter 3

Position Papers

In this chapter, we present the individual contributions of workshop participants. As mentioned previously, we had a total of 11 plenary talks and 7 talks in the work-in-progress session. We asked each speaker to provide a short abstract that summarized the ideas and opinions that this speaker aimed to present. In many cases, we received these abstracts - in these cases, the text is included with only minor edits. In many other cases, no abstract was made available to us. In those cases, we summarized the talk based on the slides and the presentation.

Within each section, the talks are sorted alphabetically. Also, note that two speakers (M. Dacier and M. Costa) decided that they do not wish their presentations to appear publicly (because of the restrictions imposed by their employers).

3.1 Workshop Plenary Talks

M. Dacier: New threats . . . What, why, who . . . ?

In his talk, Dr. Dacier was aiming to stimulate discussions. He first started off by summarizing his understanding of the workshop. In particular, he stated that the workshop, according to his opinion, aims to identify new threats that malicious actors have not started using yet. He then continued to clarify that a new threat can be a completely new threat that no one had thought about before. At the same time, it can be an old threat that had not been exploited yet.

Some open questions that Dr. Dacier put forward were:

- Do threats all rely on the same underlying technologies, and can we handle them all using the same strategies/tools/tactics?
- Do we first need to identify the various ecosystems and study them independently from each other?

- Are these various ecosystems likely to interact, cooperate together so that, by studying a (small ?) subset of them, we could indirectly learn about the others?

When looking at threats, it is important to analyze them scientifically. A scientific method consists of principles and procedures for the systematic pursuit of knowledge involving the recognition and formulation of a problem, the collection of data through observation and experiment, and the formulation and testing of hypotheses. Hence, in order to understand new threats, we need to look at the existing data.

In the past, the Symantec Internet Security Threat Report has covered 6 month periods from January 1st to June 30th and July 1st to December 31st. This report covers full year periods for the first time. All trending metrics in the report compare the full year 2007 to the full year 2008. In his talk, Dr. Dacier presented a small subset of the findings.

According to the report, attackers are now increasingly targeting end users by compromising high-traffic, trusted websites. Also, they are moving their operations to regions with emerging Internet infrastructures and, in some instances, developing and maintaining their own service provisioning. Cross-functional industry cooperation in the security community is becoming imperative.

Dr. Dacier reported that attackers locate and compromise a high-traffic site through a vulnerability specific to the site, or in a Web application it hosts. Once the site is compromised, attackers modify pages so malicious content is served to visitors.

In 2008, Symantec blocked an average of more than 245 million attempted malicious code attacks worldwide each month. Over 60% of Symantecs malicious code signatures were created in 2008. Over 90% of threats discovered in 2008 are threats to confidential information.

In 2008, Symantec identified 15,197 distinct new bot C&C servers, of which 43% were over IRC channels and 57% over HTTP. Interestingly, the United States was the country most frequently targeted by denial-of-service attacks in 2008, accounting for 51% of the worldwide total. At the same time, the United States was the top country of origin for Web-based attacks in 2008 accounting for 38% of the worldwide total.

In 2008, 95% of attacked vulnerabilities were client-side vulnerabilities and 5 percent were server-side vulnerabilities. Hence, Mr. Dacier believes that the largest threat currently are client-side attacks against end-users.

M. Behringer: Complexity and Security in ICT Systems

Modern ICT infrastructure is becoming increasingly complex, measured on a number of different parameters. Security on the other hand has been seen as inversely

proportional to complexity. If these statements hold true, then the security of ICT infrastructure is decreasing. New ways of managing complexity are needed to maintain or even increase the level of security for a system. Michael's talk discussed various ways how complexity in a system is changing, and what the possible effects on security are.

Complexity of a system can be measured intuitively by a number of factors:

- The size of the system (number of components)
- The number of communications channels per component
- The state per component
- The human interface

These parameters are not independent, but correlated: For example, if there are more components in a system, the state per component is likely to be greater, since more components need to be tracked.

Security is also linked to the above factors, but inversely: As the number of components and communications channels increases, it becomes more difficult for humans to understand all possible connections, and to design appropriate security measures. This can be hidden to some extent by a simple user interface, which limits the actions a user can execute, and consequently the security exposure of the system; however, internally the system may remain prone to security issues related to complexity, such as misconfigurations, or other operational mistakes.

Traditionally, the security approach has been to limit the degrees of freedom of a system, and thus its complexity. For example, in a secure building the number of doors to a secure zone is strictly limited, and only few employees have access. In a network, a firewall limits which component can communicate with which other component, using only a pre-determined set of communications channels. The downside to this approach is that as systems become more complex, the security policy also becomes more complex.

For each system there is a threshold of complexity up to which the system behaviour can be modeled. Above this threshold the behaviour becomes non-deterministic, and can only be simulated in an approximate way. There is a general perception that it is generally desirable to have deterministic security behaviour. New models may be required to model security on an approximate level.

Approximate security parameters already exist, for example, reputation. In this model, a client or server is not confirmed binary as "secure" or not, but one or several third parties provide subjective reputational values, which can be used locally to make a security decision. If this type of non-binary, subjective decision processes is increasingly used to control increasing complexity, then new models on controlling security of large networks need to be developed. As ICT systems become increasingly complex, it is becoming harder to control their security. New security approaches are needed, which are able to evaluate non-binary decision

values. Implementing such non-deterministic security decision models across large infrastructures is little known today and requires more research.

J.Bonneau:
Unique Security Challenges for Online Social Networks

Once a niche application for students, social networking sites have recently exploded in mainstream popularity. The largest have become complex systems with hundreds of millions of users, billions of photos, and thousands of third-party applications. The oft-derided buzzword "Web 2.0" has become prophetic as SNSs repeat the growing pains of the larger Internet. SNS operators have re-implemented existing protocols such as email, instant messaging, RSS, and OpenID within their own walls. They have implemented their own markup languages and spawned an industry of third-party software developers. Not surprisingly, SNSs are continually criticised for their perceived insecurity. The list of threats is well-known: phishing, spam, cross-site scripting, malware, data and identity theft. This talk, however, focused on what is different and argued that many of these problems are fundamentally more challenging in the SNS environment.

- Easy Social Engineering: The existence of easy access to the social graph makes many scams more effective. Phishing emails are orders of magnitude more effective from friends than from strangers. 419 scams have also become common, where a compromised account is used to request an emergency wire transfer from a "friend."
- Personal Data: Privacy concerns are intensified by the highly personal nature of uploaded information. Encouraged by SNS operators, younger users view their profile as a private space and upload highly sensitive data. Because social networks require sharing to be useful, it has proved difficult to design usable access controls. There are also many gray areas for content produced collaboratively.
- Data Centralisation: Network effects predict that a natural monopoly should arise for general-purpose SNSs, and indications are that Facebook is becoming dominant. The centralised architecture of SNSs places all user data in operator-controlled silos. This data is attractive to many third parties because it is easy to access, complete, and consistently formatted. It also contains much information that is not available elsewhere, in particular the social graph. SNS operators have retained broad legal rights to use this data however they see fit.
- Economics: The business model for SNS operators remains undefined. Most proposed revenue streams involve compromising user privacy to some extent. There are also serious questions about liability for privacy violations between the SNS operator and third-party developers.

D. Brumley:
Security that helps attackers

A defining characteristic of security research is the presence of attackers. Normally we try to develop techniques that defend against attackers. But what do we do when security research, techniques, and best practices that are intended to protect systems can also potentially benefit attackers?

First, we consider software patches. Software patches are typically assumed to help security. After all, each time a vulnerable user installs a patch, there must be one less vulnerable program in the world, so security must have improved.

We show that patches can also help attackers. In particular, we introduce the delayed patch attack. In the delayed patch attack, those who first receive a patch have a security advantage. The security advantage arises from the intuition that a new patch reveals some information about the bug being fixed, and attacker can use that information to create exploits. We demonstrate the delayed patch attack against Microsoft Windows updates. We do not require source code, and can create exploits often in a matter of minutes. We discuss the pros and cons of several possible defenses.

Second, we consider the problem of filtering out exploits as done with intrusion prevention/detection systems. We show that accurate input filters (i.e., signatures in an IDS system) also leak enough information that an attacker can use them to create exploits.

Overall, this talk focuses on the general problem of securely distributing information about vulnerabilities. We also touch on the irony that the techniques we employ were originally intended to verify software as safe.

M. Costa:
Attacks on Extensions of Commodity Operating Systems

Operating system extensions such as device drivers, file systems, and network protocols, allow users to customize systems for specific uses. New extensions for commodity operating systems appear at a rate of many thousands per year. These extensions have a defect rate 3 to 7 times higher than the operating system kernel. This poses a major security challenge because extensions are fully trusted and share the address space of the kernel.

W. Lee:
Automatic Reverse Engineering of Malware Emulators

Malware authors obfuscating their code recently began producing emulated malware. They convert native malware binaries into bytecode programs using random-generated instruction sets, paired with a native binary emulator that interprets the

instruction set. No existing malware analysis technique can reliably defeat this obfuscation technique. In his talk, Wenke presented the first work in automatic reverse engineering of malware emulators. The proposed algorithms are based on dynamic analysis. More precisely, one executes the emulated malware in a protected environment and records the entire x86 instruction trace generated by the emulator. Dynamic data-flow and taint analysis over the trace enables identification of data buffers containing the bytecode program and extraction of syntactic and semantic information about the bytecode instruction set. With these analysis outputs, Wenke and his team is able to generate data structures, like a control-flow graph (CFG), that provides the foundation for subsequent malware analysis. This was implemented in a proof-of-concept system called Rotalume. The system was evaluated using both legitimate programs and malware emulated by VMProtect and Code Virtualizer. The results show that Rotalume accurately reconstructed the original software's execution paths from its representation as unknown bytecode.

D. Gollmann: Scalable Development of Secure Software

The Domain Name System is a distributed directory system for retrieving information about hosts. It is increasingly used for access control in web applications. For example, host names are used by browsers to enforce same origin policies on cookies or on connections requested by applets; SSL relies on certificates where typically a host name appears as the distinguished name. These security mechanisms can be compromised by subverting the Domain Name System. Given that the Internet has become a critical infrastructure because of the applications deployed on it, given that many applications make use of sensitive data, securing those applications is an important task. In the context of access control and DNS, the fundamental question is then whether to make "DNS more secure," i.e. accept the fact that today DNS is asked to do more than it was originally designed for, or whether the current and emerging access control problems, e.g., in mash-ups of web applications, should be approached from first principles and alternatives to access control based on host names should be developed. In this way, we would separate the security of the infrastructure provided by web applications from the infrastructure used to facilitate communication in the Web.

P. van Rossum: Wirelessly Pickpocketing a Mifare Classic Card

The Mifare Classic is the most widely used contactless smartcard on the market. The stream cipher CRYPTO1 used by the Classic has recently been reverse engineered and serious attacks have been proposed. The most serious of them retrieves a secret key in under a second. In order to clone a card, previously proposed at-

tacks require that the adversary either has access to an eavesdropped communication session or executes a message-by-message man-in-the-middle attack between the victim and a legitimate reader. Although this is already disastrous from a cryptographic point of view, system integrators maintain that these attacks cannot be performed undetected.

In the talk, Dr. Rossum proposed four attacks that can be executed by an adversary having only wireless access to just a card (and not to a legitimate reader). The most serious of them recovers a secret key in less than a second on ordinary hardware. Besides the cryptographic weaknesses, Dr. Rossum showed other weaknesses in the protocol stack. A vulnerability in the computation of parity bits allows an adversary to establish a side channel. Another vulnerability regarding nested authentications provides enough plaintext for a speedy known-plaintext attack.

T. Jaeger: A Case for Building Integrity-Verified Communication Channels

A variety of distributed computing platforms are now prevalent, including web farms and service-oriented architectures, and several others are emerging, such as cloud computing and smart metering systems. The correct operation of these distributed systems depends on the correct operation of their constituent members. However, current approaches do not enable clients of such systems nor other members to validate correct functioning of all members. Justification of integrity is either implicit (e.g., trust that the bank administers their web site correctly) or fine-grained and limited to a single machine (e.g., based on remote attestation). Trent proposes that integrity-verification of a distributed system be obtained as part of creating a secure communication channel with that system. That is, he proposes to extend the authenticated, secure communication channels provided via SSL/TLS to integrity-verified, authenticated, and secure communication channels (integrity-verified channels). In his talk, Trent motivated the need for such a feature in distributed systems, outlined technical challenges, and described recent results showing how they enable the construction of integrity verified channels. In particular, he detailed the design and implementation of integrity monitors, VM-based services to mediate all integrity-relevant operations, and asynchronous attestation, a mechanism that enables high performance generation of integrity proofs for communications. Trent's team has built these mechanisms on Xen and SELinux systems, and in his talk, he showed his evaluation results.

C. Kreibich: Technical and Sociological Infiltration of the Underground Economy: Possibilities and Issues

Spam-based marketing is a curious beast. We all receive the advertisements "Excellent hardness is easy!," but few of us have encountered a person who admits to following through on this offer and making a purchase. And yet, the relentlessness by which such spam continually clogs Internet inboxes, despite years of energetic deployment of anti-spam technology, provides undeniable testament that spammers find their campaigns profitable. Someone is clearly buying. But how many, how often, and how much?

Unravelling such questions is essential for understanding the economic support for spam and hence where any structural weaknesses may lie. Unfortunately, spammers do not file quarterly financial reports, and the underground nature of their activities makes third party data gathering a challenge at best. Absent an empirical foundation, defenders are often left to speculate as to how successful spam campaigns are and to what degree they are profitable. For example, IBM's Joshua Corman was widely quoted as claiming that spam sent by the Storm worm alone was generating millions and millions of dollars every day. While this claim could in fact be true, we are unaware of any public data or methodology capable of confirming or refuting it.

The key problem is our limited visibility into the three basic parameters of the spam value proposition: the cost to send spam, offset by the conversion rate (probability that an e-mail sent will ultimately yield a sale), and the marginal pro it per sale. The first and last of these are self-contained and can at least be estimated based on the costs charged by third-party spam senders and through the pricing and gross margins offered by various Internet marketing affiliate programs. However, the conversion rate depends fundamentally on group actions on what hundreds of millions of Internet users do when confronted with a new piece of spam and is much harder to obtain. While a range of anecdotal numbers exist, we are unaware of any well-documented measurement of the spam conversion rate.

In part, this problem is methodological. There are no apparent methods for indirectly measuring spam conversion. Thus, the only obvious way to extract this data is to build an e-commerce site, market it via spam, and then record the number of sales. Moreover, to capture the spammers experience with full fidelity, such a study must also mimic their use of illicit botnets for distributing e-mail and proxying user responses. In effect, the best way to measure spam is to be a spammer.

In his talk, Christian presented a study they have conducted, though sidestepping the obvious legal and ethical problems associated with sending spam. Critically, this study makes use of an existing spamming botnet. By infiltrating its command and control infrastructure parasitically, they convinced it to modify a subset of the spam it already sends, thereby directing any interested recipients to servers under our control, rather than those belonging to the spammer. In turn, our

servers presented Web sites mimicking those actually hosted by the spammer, but defanged to remove functionality that would compromise the victims system or receive sensitive personal information such as name, address or credit card information. Using this methodology, Christian and his colleagues have documented three spam campaigns comprising over 469 million e-mails. They identified how much of this spam is successfully delivered, how much is filtered by popular anti-spam solutions, and, most importantly, how many users click-through to the site being advertised (response rate) and how many of those progress to a sale or infection (conversion rate).

M. Monga:
Coping with a mash-up of threats in web applications

The Web is no more a static repository of data. Increasingly, web sites offer full-blown applications to their visitors, who access them through a web browser. The first generation of web applications was server-based (the client just presented the server's static output), but now the client's contribution to the application logic is often quite relevant: the server provides data and computations that are eventually executed by a client-side interpreter embedded into the browser (e.g., a Javascript virtual machine). This kind of rich Internet applications are indeed distributed systems in which the client and the server dialogues by using HTTP requests and responses. This strategy has several advantages both for the service provider, since the computational load is partially devolved to the client, and for the end user, who experiences a greater interactivity. This strategy has enabled the spread of web applications as complex as Google Docs or OpenGoo, fully featured office suites that can be used through standard web browsers. However, web applications may expose the user to severe security risks. Although the client code executes in a sand-box, the user has currently no standard means to monitor and/or control what her/his browser will do. In fact, the sand-box just protects the integrity of the system outside the browser, but the confidentiality of the information manipulated by the browser itself (cookies, browser variables, input data, etc.) and the availability of the browser's services are at risk. The first danger is the most critical: the application could (maliciously or erroneously) disclose sensitive data to the server. The problem emerges in its maximum criticality in the so called mash-ups, i.e., web applications that combine data and functionalities from more than one source into a single integrated tool: the user could be unaware of the information exchanges among the involved actors behind the scene. Let us imagine a user who is using a service (provided by X) that mixes hotel booking (provided by A) with a city map service (provided by B): when she/he discloses her/his credit card number to the application (A needs it to perform the transaction requested by X), she/he would like to know if the sensitive data is given also to B, that apparently has no reason to access it; in fact, the implied trust relationships are rather complex. While the issue

could seem similar to the traditional problem of executing a potentially dangerous application on a user machine, it presents the following peculiarities:

- the code is often dynamically generated (by the server) and asynchronously executed (by the client): this makes static analysis approaches difficult to use;

- the application evolves very rapidly and its development might span multiple organizations: this makes unpractical (or even unfeasible when unrelated actors are involved in a new mash-up) solutions based on certification authorities or on the provider reputation.

It is worth noting that browsers normally enforce a same origin policy which states that scripts loaded from a given domain (sometimes together with the port number) cannot access data belonging to any other domain. This, however, makes things even worse: in fact, in order to overcome this limitation, mash-ups present themselves to the final user as a single-origin rich Internet application, while the mashing up of code and data from different services is hidden behind a web application proxy. Thus, the problem is to understand what the user can do to get some assurance that an application made of two parts, of which one potentially controlled and monitored by the user, does not expose her/him to unexpected security threats.

3.2 Workshop Work-In-Progress Talks

V. Ganapathy: Analyzing information flow in Javascript-based browser extensions

In his work-in-progress talk, Dr. Ganapathy described a technique for analyzing information flow in Javascript-based Browser Extensions. Browser extensions are available for all major browsers in the form of plugins, browser-helper-objects and add-ons. The focus of Dr. Ganapathy's work are Javascript-based extensions (JSEs) such as GreaseMonkey, Firebug, and NoScript as JSEs affect browser security. The sandboxing of a Javascript program in a JSE is not adequate. The Javascript, unfortunately, is not constrained by the same-origin principle. JSEs can access cookies, browsing history, and the location bar. Also, malicious web sites can misuse privileges of JSE to violate confidentiality and integrity. Dr. Ganapathy and his team have modified Firefox's Javascript interpreter to track information flow labels. Their tool helps to protect users against security violations as supposedly benign JSEs often have suspicious information flows that require white-listing.

J. Luna:
Fighting financial e-fraud: BDCT's RAFFI and eCrime projects

The financial sectors in Spain (and Catalonia, in particular) traditionally have been a pioneer in the incorporation of Information Technologies and Communications (ITC) in their business models, so its development is strongly linked to this sector. Today, most prestigious financial institutions have adopted online banking as a communication and management channel for clients, offering them the ability to perform virtually all the transactions that can take place in the real world. For some institutions, these virtual operations represent important savings. However, online banking also represents an attractive scenario for a lot of e-criminals, this being one of the main barriers to increase its use.

Financial institutions have the obligation – and desire – to safeguard the interests of their clients and offer online services, so even new channels (i.e., mobile phones and Digital TVs) must adopt measures to maximize overall security.

RAFFI is a 2-year long joint project that aims to design new algorithms and tools to prevent, detect, and mitigate the attacks that might target the various channels used for electronic banking, with a special focus on mobile devices, Web, and the Digital TV.

The leadership of this project comes from one of the biggest e-banks in Spain, thus allowing RAFFI's research and outcomes to focus toward the real needs of the financial sector. Involved partners (4 small and medium-sized businesses, 3 universities and 1 technological center) will increase their degree of knowledge and strengthen their competitive position.

V. Kisimov:
Security threats in banking systems

The presentation is made as a result of the analysis of the possible security threats that exist in banking ICT systems. Generally, the set of security threats is an open set, which extends every day. In each moment, new threats are created. Independently of this, the current presentation has an aim to classify the threat types and to establish their role in the different areas (domains) of the banking ICT system.

As a starting point to the presentation, the threat action-protection lifecycle is established, explaining the cycle from the creation of a threat to its mitigation and possible protection. The lifecycle consists of 6 steps in which threats affect the banking system. The lifecycle can be applied as an iterative process, used to identify a threat, find business exposure and decrease its effect. The threat can act in different places in the banking ICT system, and for this reason, the threat identification has to be done on macro business level - risk management, asset management, and security exposure management. With proper safeguarding measures, the threat

effect can be mitigated or stopped. All these steps are interconnected in the threat action-protection lifecycle.

One of the ways to decrease the threat effect is implementing risk management. Via risk management, the places where threats can be applied as well as their effect to the banking systems can be reduced seriously. The presentation showed a specific methodology for risk management, created by the author, through which different solution of security architecture can decrease the place and effect of security threats. After applying such a methodology, then the threat analysis has to be applied, which means there will be less places where threats can act and less effect of them can be expected.

A. Hainer: Security assessment by communities: fact or fiction?

In his talk, Andreas Hainer discussed if it is possible to do security assessment by communities. He first talked about the current security attitudes of users. It seems to be evident that users largely trust web sites and applications such as YouTube, Firefox, and Google. He then went on to talk about the increasing importance of social networking sites. He briefly presented a prototype of a software for the Nokia N810 where the feedback by the community of a user is taken into consideration before the software is installed.

J. Bonneau: Photo Sharing and Content Delivery Networks

In this short presentation, Joseph Bonneau talked about photo sharing privacy problems with respect to content delivery networks. In a study that he conducted, he created a photograph and deleted it from web sites such as Facebook, Flickr and MySpace. As these sites increasingly use external content delivery networks, the aim of the study was to find out how long it took the content providers to really delete a photograph that the user had decided to delete. In some cases, it took up to 6 days.

Also, Mr. Bonneau discovered that some of the algorithms these companies were using to generate unique IDs were not very random. Hence, it was sometimes possible to remotely access the photographs of registered users.

A. Partida: IT security incidents database: A proposal

Mr. Partida was narrating about his efforts on building a security incident database. This database should cover the details (both technical and business impact) of real

cyber security incidents. Mr. Partida was emphasizing the advantages of such a database with real cases and data, in particular, when attempting to motivate support and funding for security improvements. As he put it: “Nobody pays attention when one uses statistics that describe the increase in cyber security incidents. However, when providing concrete examples of what can happen and what damage can be done, people are much more inclined to listen.” Of course, contributors to this database can decide to which extent they are willing to share information, and the confidentiality of information has the highest priority. Alberto Partida then called for potential contributors to contact him.

O. H. Longva: Activities at SINTEF

In his talk, Mr. Longva talked about three ongoing projects at SINTEF. In the Semantic Information Security in Integrated Operations (SIRCUS) project, the aim is to develop a security and trust framework for the Semantic Web to enable the implementation of secure technologies for communicating organisations supporting Integrated Operations offshore.

In the Intelligent Access Control and Misuse Detection for Healthcare Information Systems (iDetect) project, the aim is to have improved access control and misuse detection in healthcare systems that can be achieved by utilizing traces of user actions recorded in existing audit logs.

In the Metrics for information security incident response (TRICSI) project, SINTEF is working on developing a set of indicators for incident management and methods for measuring these indicators automatically.

Chapter 4

Conclusions

The consortium feels that the workshop was a great success. For one, the planned number of participants was significantly exceeded (103 participants compared to a success threshold of 60). Moreover, we managed to attract a high number of well-known researchers from around the world. In fact, two of the invited talks for the workshop received the best paper awards at the highly selective IEEE Security and Privacy conference that was held a week later. In addition, we also received positive feedback that demonstrated the need for an initiative such as FORWARD .

Based on the concrete feedback we received during the workshop and, in particular, the discussion sessions, each working group can be confident that all relevant threats have been covered. Moreover, an initial ranking and prioritization step was carried out. In the separate working group meetings, we witnessed lively discussions by the participants. The general feeling at first was that it is not easy to rank the threats. However, by guiding and channelling the discussion, we rated the threats based on issues such as impact and cost. This outcome will serve us very well during the third and final project phase, in which the importance and the inter-dependencies among threats will be analyzed.

With the second workshop, we have created a tremendous visibility for the FORWARD project, not only in Europe but also among the researchers in the US and other countries. This is a momentum that we can continue to build upon in the future, and a community building effort that the European systems security landscape will considerably profit from.

Chapter 5

List of Participants

Adam Kozakiewicz, NASK
Alberto Partida
Aljosa Pasic, Atos Origin
Andrea Lanzi, Institute Eurecom
Andreas Heiner, Nokia Research Center
Angelos Keromytis, Columbia University and Symantec Research Labs Europe
Angelos Stavrou, George Mason University
Asia Slowinska, Vrije Universiteit Amsterdam
Bart Jacobs, Institute for Computing and Information Sciences, Radboud University Nijmegen
Christian Kreibich, ICSI
Christian Platzer, Technical University of Vienna
Christopher Kruegel, TU Vienna
Claude castelluccia, INRIA
Corrado Leita, Symantec Research Europe
Cedric Blancher, EADS Innovation Works
Damiano Bolzoni, University of Twente
Daniel Keim, University of Konstanz
Daniel Haglund, Swedish Civil Contingencies Agency
Darren Mutz, Virgin Charter
David Barroso, S21sec
David Brumley, Carnegie Mellon University
Davide Balzarotti, Institute Eurecom
Dieter Gollmann, Hamburg University of Technology
Dimitrina Polimirova, National Laboratory of Computer Virology - BAS, Research associate
Diomidis Spinellis, Athens University of Economics and Business
Dirk Westhoff, NEC Europe Ltd.
Edita Djambazova, IPP - BAS
Emma Crowe

Engin Kirda, Institute Eurecom
Erik Blass, Institute Eurecom
Erland Jonsson, CHALMERS
Eugene Nickolov, National Laboratory of Computer Virology - BAS, CEO
Evangelos Markatos, FORTH-ICS
Fabian Monrose, University of North Carolina, Chapel Hill
Fang Chen, Chalmers University of Technology
Frank Fransen, TNO ICT
Georg Kremsner, IKARUS Security Software GmbH
Georgios Portokalidis, VU Amsterdam
Gerhard Paass, Fraunhofer IAIS
Guevara Noubir, Northeastern University
Guy Bunker, Symantec Corporation
Herve Debar, France Telecom RD
Holger Dreger, Siemens AG, Corporate Technology
Hong-Linh Truong, Vienna University of Technology
JAMES CLARKE, Waterford Institute of Technology
Jan Kohlrausch, DFN-CERT
Jesus Luna, Barcelona Digital
John Ioannidis, Google
Jonathan Smith, U Penn
Jonathon Giffin, Georgia Institute of Technology
Joseph Bonneau, University of Cambridge
Kiril Dimitrov, IPP-BAS
Kostas Anagnostakis, A*STAR/Niometrics
Levente Buttyan, Budapest University of Technology and Economics
Leyla Bilge, Institute Eurecom
Luben Boyanov, IPP-BAS
Magnus Almgren, Chalmers
Manolis Stamatogiannakis, FORTH/ICS
Manuel Costa, Microsoft Research
Marc Dacier, Symantec Research Europe
Marc Heuse, Baseline Security, DIMVA, GI eV, ...
Marco Balduzzi, Institute Eurecom
Mario Miladinov, Information Services
Martin Koyabe, BT
Massimo Ciscato, European Commission
Mattia Monga, Universita degli Studi di Milano
Melek Onen, Institute Eurecom
Michael Bailey, University of Michigan
Michael Behringer, Cisco Systems
Mihai Christodorescu, IBM Research
Moti Yung, Google (and Columbia U.)
Neeraj Suri, TU Darmstadt

Odd Helge Longva, SINTEF
Oznur Ozkasap, Koc University
Paolo Milani, TU Vienna
Peter van Rossum, Radboud University Nijmegen
Philip Homburg, Vrije Universiteit
Philipp Trinius, University of Mannheim
Piotr Kijewski, NASK
Refik Molva, Institute Eurecom
Richard Kemmerer, UCSB
Roberto Cannata, ENEA
Roberto Di Pietro, EUROJUST (Data Protection Service)
Rogier Spoor, SURFnet
Ruth Fochtner, TU Vienna
Sathya Rao, Telscom
Simin NADJM-TEHRANI, Linkoping university
Sotiris Ioannidis, FORTH-ICS
Spiros Antonatos, ICS-FORTH
Steven Bellovin, Columbia University
Thorsten Holz, University of Mannheim
Thorsten Strufe, TU Darmstadt
Trent Jaeger, Penn State University
Urko Zurutuza, Mondragon University
Valentin Kisimov, Bulgaria
Vassilis Prevelakis, Head of Research, AEGIS RESEARCH
Vinod Ganapathy, Rutgers University
Virgil Gligor, Carnegie Mellon University
Wendy M Grossman, freelance
Wenke Lee, Georgia Institute of Technology
Wietse Venema, IBM Research
William Robertson, UC Santa Barbara
Wolfgang Trexler, Bank Austria