Virtualized Security at the Network Edge: A User-Centric Approach

Diego Montero, Marcelo Yannuzzi, Adrian Shaw, Ludovic Jacquin, Antonio Pastor, René Serral-Gracià, Antonio Lioy, Fulvio Risso, Cataldo Basile, Roberto Sassu, Mario Nemirovsky, Francesco Ciaccia, Michael Georgiades, Savvas Charalambides, Jarkko Kuusijärvi, and Francesca Bosco

ABSTRACT

The current device-centric protection model against security threats has serious limitations. On one hand, the proliferation of user terminals such as smartphones, tablets, notebooks, smart TVs, game consoles, and desktop computers makes it extremely difficult to achieve the same level of protection regardless of the device used. On the other hand, when various users share devices (e.g., parents and kids using the same devices at home), the setup of distinct security profiles, policies, and protection rules for the different users of a terminal is far from trivial. In light of this, this article advocates for a paradigm shift in user protection. In our model, protection is decoupled from users' terminals, and it is provided by the access network through a trusted virtual domain. Each trusted virtual domain provides unified and homogeneous security for a single user irrespective of the terminal employed. We describe a user-centric model where non-technically savvy users can define their own profiles and protection rules in an intuitive way. We show that our model can harness the virtualization power offered by next-generation access networks, especially from network functions virtualization in the points of presence at the edge of telecom operators. We also analyze the distinctive features of our model, and the challenges faced based on the experience gained in the development of a proof of concept.

Diego Montero, Marcelo Yannuzzi, and René Serral-Gracià are with Technical University of Catalonia (UPC).

Adrian Shaw and Ludovic Jacquin are with Hewlett-Packard Laboratories, United Kingdom.

Antonio Pastor is with Telefónica I+D.

Antonio Lioy, Fulvio Risso, Cataldo Basile, and Roberto Sassu are with Politecnico di Torino.

Francesco Ciaccia is with Barcelona Supercomputing Center (BSC).

Mario Nemirovsky is with ICREA Researcher Professor at BSC.

Michael Georgiades and Savvas Charalambides are with PrimeTel PLC.

Jarkko Kuusijärvi is with VTT Technical Research Centre of Finland Ltd.

Francesca Bosco is with United Nations Interregional Crime and justice Research Institute.

INTRODUCTION

The protection of users' terminals against Internet threats is largely dominated by a device-centric model. This basically consists of installing a set of security applications on each terminal, such as anti-virus software and a personal firewall. An average user nowadays has multiple terminals, including a smartphone, a smart TV, and a notebook, and in many cases also a tablet, a desktop computer, and even a game console. These devices usually have different architectures (e.g., Intel or ARM) as well as different capabilities and operating systems (e.g., Android, Windows, or Linux), so the appropriate protection tools may not be available for all platforms. As a result, the most common practice is to install different security applications on the various terminals — or simply rely on the default protection means provided by the operating systems. Let us assume for a moment that users would like to have the same security policy and exactly the same protection level enforced on all of their devices. In the context of this article, we will call this the "uniform security aim." To achieve this goal, the user typically needs to understand the configuration details of each device, which typically involves the setup of different security applications on different platforms. For non-technically savvy people, this turns out to be an impossible hurdle to overcome. As a result, most Internet users suffer from wide variations in their protection levels, and this problem is exacerbated as the number of devices per user grows.

In this article, we propose a paradigm shift from device-centric protection to a user-centric model. The latter specifically addresses the two main drawbacks of the former: the need for dissimilar installations of security applications in different devices due to their different platforms, and the problem of non-uniform protection due to the difficulties in the configurations needed.

To cope with the first problem, we propose a model in which the protection and security policies are now unified and remain homogeneous for each user, independent of the terminal used. This is achieved by means of a user-specific trusted virtual domain (TVD), which is dynamically instantiated at a secure place in the network edge. As we shall show, the TVD can be instantiated either on the user's side (e.g., on a home gateway) or on the provider's side (e.g., on a next-generation broadband access server handling the user's connections).

To cope with the second problem identified above, we propose a user-defined security model that aims at ease of use by design. We discuss the importance of exposing the selection of high-level protection policies to the average user, and the necessity to enforce the configurations required transparently to the latter. This simple strategy detaches the definition of the protection policies from their corresponding configurations, thus allowing tailored protection even by non-

0163-6804/15/$25.00 © 2015 IEEE

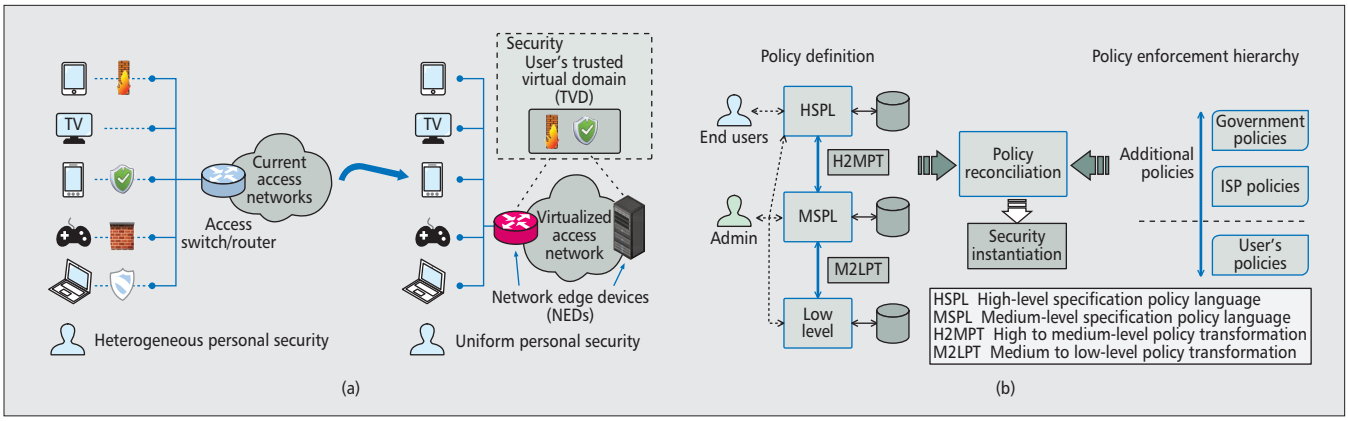

Figure 1. The two main objectives of our user-centric model, that is, uniform protection and ease of configuration: a) offloading security to the virtualized access network; b) policy definition and the policy enforcement hierarchy.

technically savvy users. It is worth highlighting that the virtualized security model described in this article can be applied both to residential and corporate scenarios. We describe its application in the form of a multi-tenant platform, considering the main stakeholders involved (i.e., service providers, infrastructure providers, security application developers, and users).

The remainder of the article is structured as follows. First, we outline the essentials of the paradigm proposed, including the new protection model and the security policy approach. Next, we introduce the general architecture and its main components. After that, we analyze the distinctive factors of our model, and outline some of the main conclusions that can be drawn from our prototype implementation. Finally, we conclude the article.

Toward a New Protection Paradigm

Figure 1a depicts the basic concepts, showing the evolution from device-specific security to a common security framework for all devices hosted in the access network. In our model, security applications that are commonplace today (anti-virus, firewalls, content inspection tools, etc.) shall be called personal security applications (PSAs). Observe that under the current protection model, the heterogeneity of devices and platforms requires the installation of various PSAs with similar roles and functions; actually, four PSAs are required in the example shown in Fig. 1a. Also observe that some devices may remain completely unprotected, as in the case of smart TVs.

Under our paradigm, the heterogeneous set of PSAs protecting the different devices is now moved and consolidated into a TVD. Each TVD only needs to host the minimum set of complementary PSAs required by the user (e.g., an anti-virus and a firewall in the example). A TVD is a "logical container" that is instantiated *per user*, and is composed of the following elements:

- The execution environments hosting the user's PSAs
- The required data, control, and management plane interconnectivity in order to guarantee the isolation between different users' TVDs (we delve into this later; Fig. 3).

The right side of Fig. 1a shows that a user TVD can be instantiated at either end of the access link. Indeed, as a logical container, a TVD may run entirely within a single network edge device (NED), or in a distributed way involving several NEDs. In our terminology, a NED is a device with virtualization capabilities that supports the instantiation of TVDs in a multi-tenant fashion. If the TVD is placed in a user's premises, the NED could be either an enhanced home gateway or customer premises equipment (CPE). Those devices may need additional compute, storage, and networking resources, and could be managed by the Internet service provider (ISP). If the TVD is placed in the ISP premises, as will be the case with the upcoming network functions virtualization (NFV) based access networks [1], a pool of nodes belonging to the NFV infrastructure could be the NEDs devoted to host our TVDs. Note that this second deployment strategy leverages the virtualization and processing power of commodity hardware, and the unquestionable trend toward its ubiquity at the network edge — although it does not exclude the adoption of the first deployment strategy as well. It is worth highlighting that our model has a remarkable advantage over cloud-based protection [2]. Whereas in the latter case the virtualized resources supporting the users' security are rarely on the path that would naturally be followed by user traffic, in our model, the TVD is always instantiated on the natural path. In other words, our model avoids routing detours, which would occur if the NEDs were located off the path between the user terminal and its traffic destinations (e.g., in the cloud).

As its name indicates, the TVD must be trusted, since it will execute security applications on behalf of the user on one or more nodes that are typically owned and managed by a third party. Appropriate techniques, such as remote attestations [3] or contractual agreements, must be put in place to guarantee the appropriate level of trust according to the security needs of a specific user. Also, observe that the NEDs must be secure, since they will host the applications of several users that could potentially affect each other. As we shall show, the NED must be connected with a secure channel to the user termi-

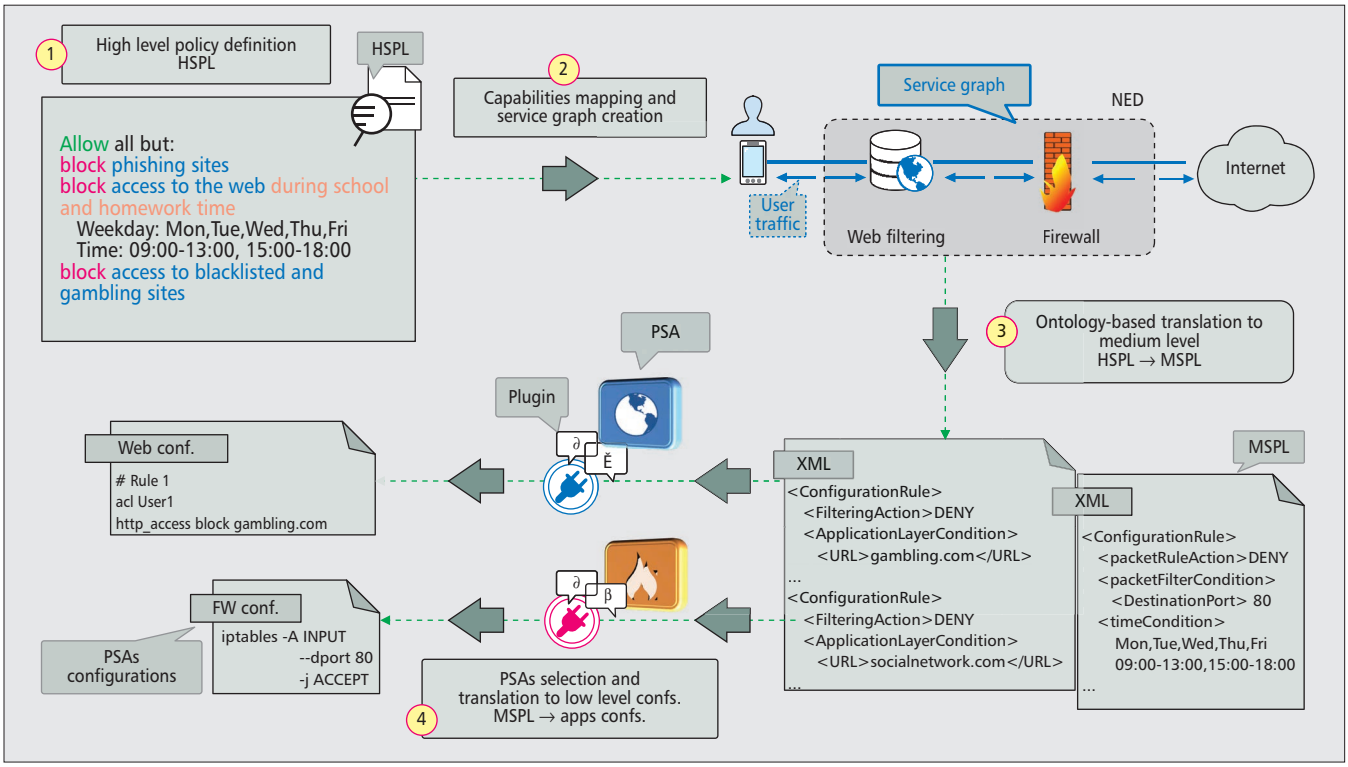

Figure 2. Example of policy definition and enforcement, going from HSPL to MSPL and then to low-level configurations.

nal, because this path may be subject to attacks that could try to bypass the security controls performed at the NED.

Each PSA within a TVD implements one (or possibly more) security controls that need to be configured according to the needs of a specific user. However, the configuration of security applications is often complex and not well understood by the majority of users. To simplify this task, we propose the model shown in Fig. 1b. The rationale behind it is that to build a real user-centric model, it is mandatory to allow users to specify their own security requirements (i.e., their *security policy*) in a straightforward way. Our design principle aims to meet the expectations of both non-technically savvy users and experts in the field, such as security administrators. For the former, the goal is to allow them to specify their security policy without needing to deal with the technicalities. For the latter, the goal is to allow them to fine-tune their policies while simplifying the configuration of the security applications under their administration.

To achieve these goals, our model is composed of three policy abstraction layers, and two translation services between them (see the left side of Fig. 1b). The first abstraction layer is supplied by the High-Level Security Policy Language (HSPL), a user-oriented authorization language suitable for expressing concepts related to user protection. HSPL allows users to express general protection requirements by means of sentences that are very close to natural language, such as "do not permit access to war content," "block my son from accessing gambling sites," or "allow email scanning." In our model, HSPL policies can be selected from a set of candidate policies that can be then customized and grouped (e.g., "block access to gaming sites" + "only during weekdays"). The policy sentences are internally mapped to a *subject-verb-object-attribute* authorization language that is currently under definition as an XACML profile [4]. For instance, the policy "block my son from accessing gambling sites" is interpreted as "block" (verb) "my son" (subject) "from accessing gambling sites" (the object). Predefined lists of subjects, verbs, and objects are made ready for the users, so they can easily compose their own sentences. Available attributes depend on the verb-object pair. Moreover, users can extend the predefined fields without being experts. The specific details of HSPL are out of the scope of this article, so for additional information the reader is referred to [5].

The lowest layer in the policy abstraction stack is what we call the "low level" in Fig. 1b, as it is the one that deals with the configuration details of the PSAs. This configuration procedure is clearly application-specific, and hence is not under our control.

With the aim of abstracting the specific configuration procedures while meeting the experts' needs, we have created an intermediate abstraction layer that allows the specification of PSA configurations using a PSA-independent format. The security policies in this abstraction layer are specified by means of the Medium-Level Security Policy Language (MSPL). The effort in the definition of the MSPL is not trivial. Indeed, depending on the heterogeneity of the different security control languages, the mappings can be arbitrarily complex. We address this complexity by means of an MSPL model that defines the main concepts (e.g., policies, rules, conditions, and actions), and is organized by *capabilities*. In

this context, capabilities are defined as basic features that can be configured to enforce a security policy (channel protection, filtering, anti-virus, parental control, etc.). Our approach also allows families of languages with similar concepts to be grouped (e.g., attributes, actions, or condition types), which can be captured by specific submodels built by analyzing several languages of controls sharing the same capability. For instance, through MSPL it is possible to write the configuration of a general packet filter or to configure the options of a general anti-virus. An illustrative example of MSPL outlining the translation from HSPL up to low-level configurations is sketched in Fig. 2.

Overall, writing policies in MSPL demands the same security awareness and level of expertise as specifying the configurations directly in the PSAs. The advantage, however, is that MSPL spares experts the burden of mastering several semantically equivalent security controls and syntaxes. Observe that PSA developers will need to provide their plug-ins jointly with their PSAs, in the form of a medium- to low-level policy translation (M2LPT) service (Fig. 2). Also note that the complexity mainly resides in the language definition, so these translators fundamentally perform syntax adaptation. Thanks to this approach, a security policy written in MSPL can be embodied by different PSAs, provided that the candidate PSAs offer the capabilities required by the user. In addition, the PSAs can be replaced without impacting the security policy specified by the user (e.g., replacing a Cisco packet filtering application by one provided by Checkpoint). For further details on MSPL, the reader is referred to [5].

As shown in Fig. 1b, the binding between HSPL and MSPL is supplied by the high- to medium-level policy translation (H2MPT) service. Different from the M2LPT translation, which is provided by the PSA developer, H2MPT represents a translation service that is natively provided by our architecture. H2MPT uses formal ontologies to provide the semantics implied by the high-level policy statements. Our ontology is based on [6], and it models the high-level concepts (subjects, objects, verbs, and attributes) as well as the medium-level concepts (rules, conditions, actions, resolution strategies) and the capabilities. The ontology also contains information on how predefined HSPL concepts are expanded into useful information for building MSPL rules. The translation process first identifies a set of applications that can enforce the security policies (e.g., a web filter and a firewall), and then generates the MSPL for the selected applications. The HSPL verb-object pairs are used to match the capabilities needed for policy enforcement, while the capabilities per se are used to determine the PSAs and their interactions.

Moreover, a meta-model defines how HSPL sentences are mapped into MSPL concepts, and how these concepts must be assembled to build valid rules. This meta-model is used by a set of enrichment modules and by a standard ontology reasoner to gather all the information needed to create MSPL policies that enforce the HSPL policy [5, 6]. Finally, an H2MPT component combines this information into MSPL policies. This translation is done transparently for nontechnically savvy users (i.e., for those users specifying their policies through HSPL). We contend that by having a high-level policy specification language, our model provides far more flexibility and expressiveness than approaches based on profiles or templates. This is because these latter basically wrap under a common name a set of low-level settings, which are basically applied for a fixed set of security controls.

In the model we conceive, the PSAs can be selected by the users themselves or by a provider. If the user only specifies the HSPL, the PSAs are automatically selected from a catalog of available applications based on the PSAs that meet the functionality required by the policies. In our model, the capabilities of a PSA are specified through a "PSA manifest." In this context, the selection may be straightforward — when only one PSA is available with the required capabilities — or it may be based on various criteria if multiple PSAs could offer those capabilities (on the PSA reputation or its cost, etc.).

Another important aspect is that according to recent studies, human mistakes are the major cause of breaches and vulnerabilities [7]. Thus, our model provides analytics that help reduce the likelihood of such mistakes. These include contradictions among policies in different PSAs, policy contradictions within a PSA, or cases leading to suboptimal performance (e.g., rules that are never matched and simply increase the processing time). Our model identifies these types of anomalies by means of state-of-the-art techniques [8]. We represent clauses as hyperrectangles so that anomalies can be detected by using geometric intersections. Anomalies are classified by evaluating geometric relations among conditions (e.g., inclusion, intersecting conditions but no one includes the other), as well as relations between actions (e.g., same action, equivalent actions, conflicting actions). The resolution is dealt with by formally modeled strategies, which cover a set of existing security control resolution mechanisms. Upon detection, we provide hints on how to resolve them and notify the effects of each decision.

Moreover, the model we envision should support multiple actors, which could simultaneously operate on the same traffic (see the right side of Fig. 1b). Each of these actors may possibly impose its potentially conflicting security policy. For instance, a user can decide the level of protection needed, but the ISP may impose other limitations in order to guarantee the integrity of its network. In turn, the government may impose additional restrictions. In the case of conflict between the different policies in the hierarchy, our approach is to automatically resolve such anomalies, and inform the user about the issue and its outcome.

In order to resolve such conflicts, a "reconciliation" [9] process is performed. The latter takes the policies of the different actors that must be reconciled, and obtains a single MSPL policy to be enforced by the user's PSAs. The core of this process is the resolution of contradictions among rules from different policies. Priorities and hierarchies are some of the simplest ways to resolve

A derived requirement posed by multi-tenancy is network isolation. The SECURED architecture must ensure the isolation of traffic among different users. More precisely, each tenant will be configured with a dedicated and private virtual network.

contradictions (i.e., rules from higher-priority policies/actors prevail), and they typically map well to contractual frameworks. However, custom reconciliation strategies can be defined. The reconciliation process copies non-conflicting rules in the reconciled policy, while each resolved contradiction generates a new rule. The latter have higher priority than the original ones, and the correct action is decided by the selected reconciliation strategy. More details on our reconciliation approach can be found in [10].

Observe that actors may decide not to disclose their policies to other actors. In that case, reconciliation strategies that require full access to the policy set are not possible. An alternative approach is to use *policy chaining*. This consists of redirecting the output of one set of PSAs in an administration domain (e.g., the user PSAs) to a set of PSAs in another domain (e.g., the ISP PSAs). The user must not necessarily own the PSAs in other domains when chaining is performed. This is useful when more sophisticated controls are required by the entities that specify the higher policies in the stack.

An Example of Policy Translation and Enforcement

To better describe our new paradigm, we present an example that illustrates the step-by-step process, starting from the definition of high-level policies up to the configurations made to guarantee their enforcement. Figure 2 depicts a simplified but complete example of the policy definition process for a non-technically savvy user. It comprises four basic steps. First, the user is requested to define its policies using HSPL. This user-oriented authorization language allows a set of general security rules to be expressed and customized by means of sentences that are very close to natural language (e.g., *block phishing sites*).

Next, the HSPL policy sentences are mapped to a subject-verb-object-attribute authorization language aiming to extract the different security capabilities required by the user (Fig. 2, step 2). As a result, a service graph is built, where the nodes represent generic applications (PSAs) capable of fulfilling the security requirements. Observe that two applications are required in the example, web filtering and a firewall. The selection of PSAs is based on the manifest provided along with each PSA, which indicates its specific capabilities. Third, by using the ontology and the service graph information, the security policies are translated into MSPL, obtaining the application-independent definition of policies requested by the user (Fig. 2, step 3). The representation of MSPL policies is stored and managed in XML format. Fourth, specific PSAs are selected satisfying the capabilities and requirements of the user. As mentioned above, the specific PSAs can be selected by either the user or the provider. For each PSA, the configurations are created using an application-specific translation plugin. These plugins convert the generic MSPL rules to application-specific configurations (Fig. 2, step 4). These configurations will be the inputs once the PSAs are instantiated and linked.

Finally, once the PSA configurations are created, an orchestration system instantiates each PSA and enforces its particular configuration, hence providing the security policies defined by the user.

The SECURED Architecture

This section introduces the envisioned architecture, which we call SECURED [5]. As explained above, SECURED provides a system where users can offload their PSAs to their nearest compatible NED. The architecture is specifically devised to be heavily multi-tenanted and flexible enough to be used in scale-out systems. From a use case point of view, it can be expanded and deployed in a variety of ways, ranging from small set-top boxes or home gateways up to deployments on a much larger scale in a distributed environment (e.g., in localized data centers at the edge of ISP networks). Our focus in this section is on the main architectural components.

General Overview

The architecture must support the dynamic allocation and instantiation of users' security. The security functionality of each user can be comprised of different PSAs in a defined arrangement through service chaining, and these PSAs can be deployed within the same physical host or in a distributed manner. As a result, two general requirements are imposed on the architecture:

- Massive multi-tenancy, which implies isolation of users, their applications, and network traffic
- A secure and verifiable infrastructure and environment, which users can trust to host their security applications

A general view of the basic architecture is depicted in Fig. 3. The figure shows a generic deployment (e.g., on an NFV POP of an ISP). It is worth noting that in simpler deployments (e.g., when the NED is a home gateway), the functionality provided by some of the systems at the top of Fig. 3 could be simplified and embedded in the NED itself, or might not be needed, such as the case of the NFV orchestrator.

Overall, the first requirement is to guarantee complete isolation between different users. In light of this, the TVD was designed as an isolated environment that will hold the security applications of a user and in turn process the user traffic. A TVD comprises one or more execution environments (EEs). An EE is a lightweight and heavily controlled environment that contains and executes one or more user PSAs,each operating on the principle of least privilege. Thus, within SECURED, two levels of isolation are defined (Fig. 3):

- The *compartmentalization layer*, which is mainly responsible for the isolation between user TVDs
- The *containment layer*, which handles isolation between PSAs within an EE

Thus, an EE could be either a compartment or containment layer, respectively.

A derived requirement posed by multi-tenancy is network isolation. The SECURED architecture must ensure the isolation of traffic among different users. More precisely, each tenant will be configured with a dedicated and private virtu-

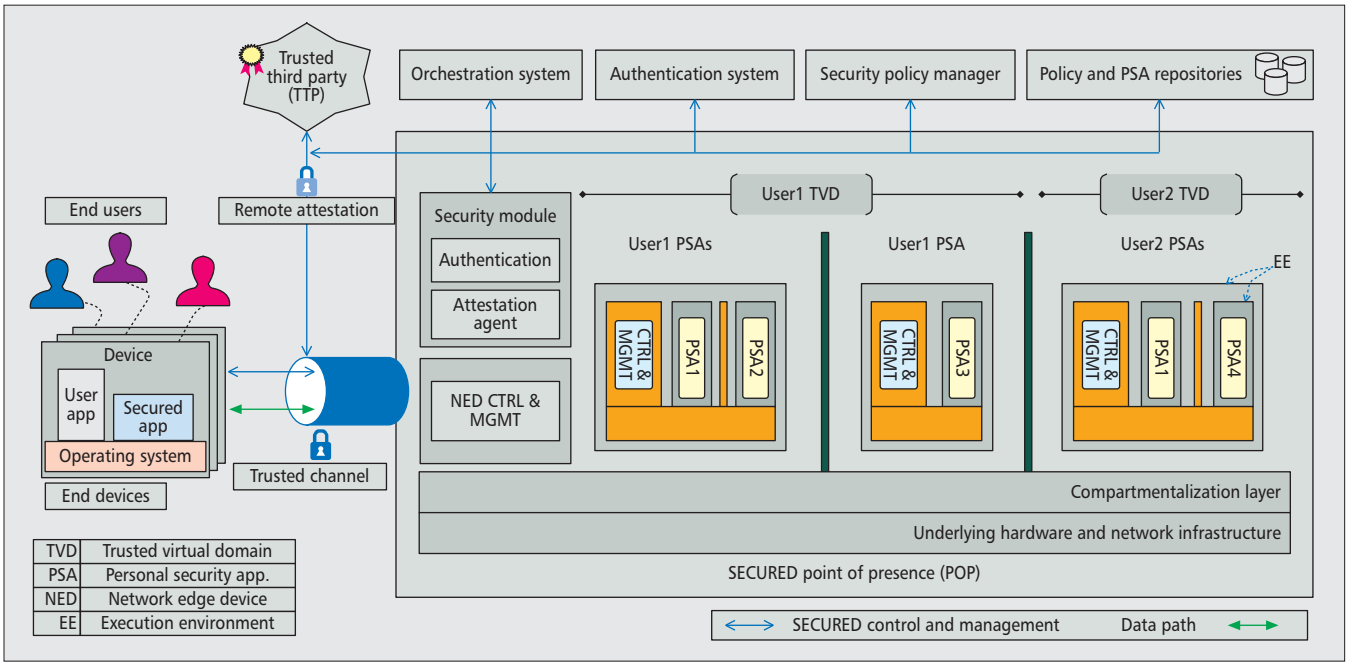

Figure 3. The basic SECURED architecture showing a multi-tenant scheme on a point of presence (POP).

al network. This network connects the different PSAs with the end user on one side and the Internet on the other. Furthermore, the architecture defines a private management network that sets up, controls, and manages the different TVDs. Both the compartmentalization and containment layers have a *control and management* component, which aims to establish separation between the technology-independent part and the implementation-dependent technology.

The second requirement is related to the establishment of trust between the end user and SECURED. This requirement is vital, since users would like to establish a certain level of trust with SECURED prior to requesting the instantiation of security applications and sending their traffic. We address this requirement by using the concept of *remote attestation* (RA). SECURED leverages trusted computing mechanisms to measure the system software upon component startup, where resulting measurement digests are held by a secure root of trust, such as a hardware device like a trusted platform module (TPM) [11]. These measurements can be cryptographically signed by the device and sent to the users whenever they send an attestation request. The process of RA poses a major challenge for SECURED, and preliminary insight on a proof-of-concept implementation is described later.

Main Components

Security Module — This module is the front-end, which is contacted during connection establishment. It comprises two elements, the *attestation agent* and the *authentication module*. Prior to authenticating, the end user first contacts SECURED in an attempt to establish a secure connection while also performing the remote attestation protocol. To this end, the SECURED system receives a challenge request to perform an attestation of its software configuration. A mutually trusted third party (TTP) system is involved in the attestation process. The TTP is responsible to keep a copy of known-good measurements, and provide a secure verification service to the user for verifying remote attestation responses. After a successful check, a secure channel is created, and the user safely sends her credentials to the *authentication module.*

Authentication System — The authentication of users is a key component of SECURED. This can be implemented either using a local (standalone) authentication system or relying on an existing external authentication infrastructure (e.g., an AAA+ system). The result of the authentication process is to obtain tokens allowing the interplay between the main components within a NED, and external subsystems such as PSA repositories. Once the user is authenticated, the instantiation of his security must be enforced.

NED Control and Management — Once the user is authenticated, this module retrieves the user policies and metadata related to the composition of the required security applications. After that, the control and management module drives the instantiation of the user TVD, including its applications and setup of the virtual network. More specifically, this module determines the resources required for the user TVD, and commands the instantiations required as well as the deployment and interconnection of the PSAs. This computation encompasses an analysis of the required compartments, containments, and virtual networks to be allocated in order to instantiate the security applications. This analysis considers the PSA requirements along with the availability of resources, and the required configuration of the network (physical and virtual). In addition, this module also manages the extension

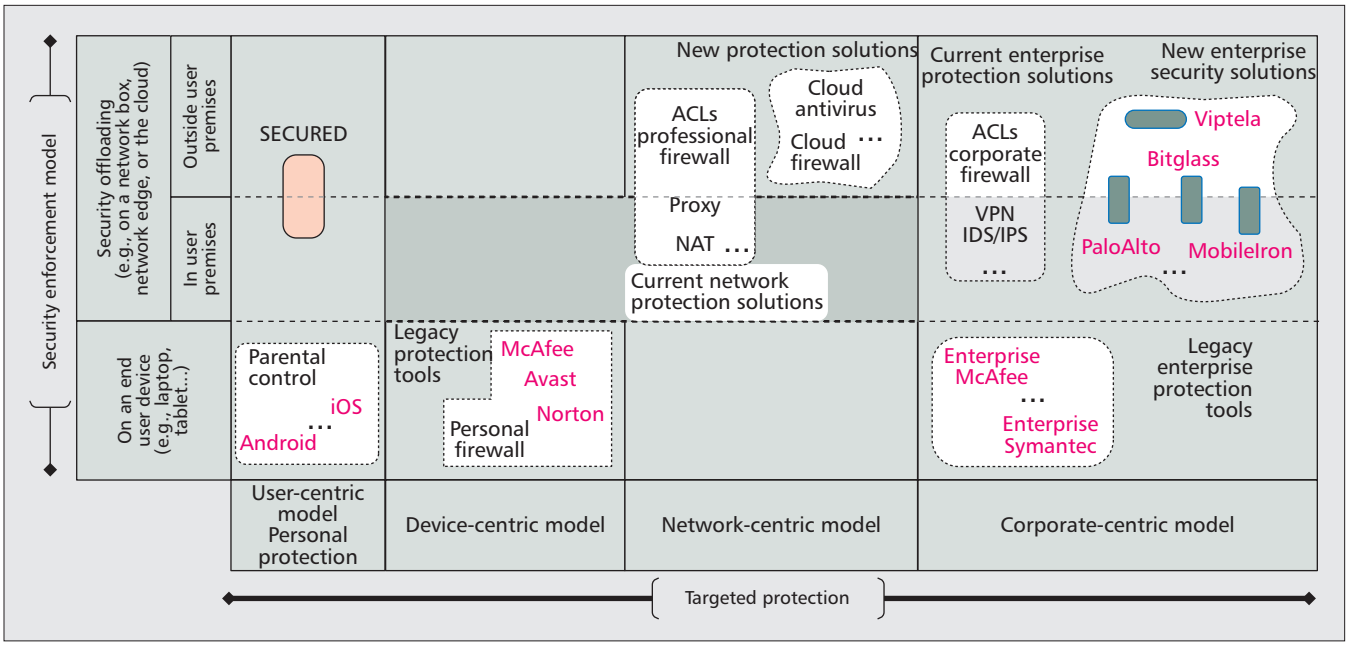

Figure 4. Positioning SECURED considering some of the most common tools as well as some of the most recent and compelling solutions in the area.

of the user data path to connect the user's device to the newly created TVD.

Orchestration System — In the case of an NFV POP, the NED control and management module will be assisted by the NFV orchestration system. However, in simpler scenarios, the former could entirely handle all the configurations required. In other words, when the NED is embodied in the home gateway of a residential user, the orchestrations needed will be handled locally without requiring any external orchestrator. In general terms, the orchestration system should deal with the instantiations and configurations in large distributed systems (e.g., an NFV POP), preferably in a "technology-agnostic" way. The technology-dependent part could be managed by the control and management module embedded in the NED. In our model, the attestation agent keeps track of the different components during the instantiation phase (i.e., compartments, containments, and PSAs), and manages the corresponding measurements in order to present an attestation proof back to the user concerning her TVD.

Security Policy Manager — This module is in charge of handling the user's policies and the reconciliation process prior to performing the configuration of the user's PSAs.

PSA Repositories — The applications are retrieved from these repositories with their respective MSPL plugins, which then need to be loaded into one or more TVD containments.

SECURED App — This is the only application that needs to be installed in a user device. Its role is basically to support the secure communications with the NED, and handle the remote attestations and their outcomes.

Overall, the architecture introduced in this section allows the dynamic creation of trusted and virtualized execution environments throughout the access network. In this framework, several actors such as users, corporate information and communications technology (ICT) managers, infrastructure providers, security service providers, and software developers can interplay and benefit from our user-centric protection model. An important remark about the proposed architecture is its alignment with the emerging NFV technology. NFV is an enabler for SECURED, and will be essential for guaranteeing its scalability.

Analysis of SECURED

The security model proposed in this article has several distinctive factors that make it unique. To show this, we position SECURED in the current spectrum of protection techniques and highlight its main differences with state-of-the-art solutions. In addition, we present and discuss our initial evaluations of a proof-of-concept implementation, with special focus on performance aspects related to the security, trust, and service verification offered by SECURED.

Positioning SECURED within the Security Panorama

The spectrum of solutions designed to counter security threats is really broad. The solutions available today can be reasonably categorized according to the table shown in Fig. 4. As can be observed, there are solutions that are focused on protecting the end-user device, while others propose different forms of security offloading. Moreover, current protection schemes can be classified based on whether they are user-centric, device-centric, network-centric, or corporate-centric. In a nutshell, Fig. 4 presents a high-level comparison of different security protection

schemes according to two general criteria: the targeted protection model and where the security is enforced.

To the best of our knowledge, SECURED is the only solution available nowadays that proposes a true user-centric model which specifically addresses the need for device-independent security. As described previously, the user-centric approach is achieved thanks to the HSPL and MSPL languages, and the H2MPT and M2LPT translation services between the three abstraction layers involved. This allows users and even experts in the field to focus on their security policies rather than on the configuration details of specific security applications. Another important aspect is that, in contrast to many of the offloading solutions available today, which are typically deployed in the cloud, our solution admits a rich variety of deployments on either edge of an access link. Cloud-based solutions provide compelling protection schemes while avoiding several of the overheads for end users (e.g., corporate customers). The downside, however, is that they require routing detours, are not really user-centric (at least not yet), do not provide essential trust means such as remote attestation, and do not support advanced features such as anomaly verification and policy reconciliation techniques. These latter two are a couple of distinctive aspects in SECURED, and therefore are the center of our assessment and analysis at this stage. We proceed to provide insight foresee based on a proof-of-concept implementation.

Remote Attestations

Trust establishment between an end user and the protection platform is a critical step toward security offloading. In our model, we use remote attestations (RAs) and verification techniques for the trust establishment process. Let us assume the following scenario: A user connects through an insecure channel and requests protection from SECURED. Prior to starting traffic exchange, the user is requested to create a trusted channel toward a NED. A trusted channel is an instance of a secure channel (e.g., a virtual private network, VPN), where the endpoints are attested before any data exchange. In SECURED, the trusted channel protects users against a potentially compromised NED. However, enabling these security countermeasures introduces overhead. On one hand, users may experience delay during the establishment of the connection with the NED. This is due to the integrity check needed, which is issued only once per user during the connection. Likewise, administrators may face scalability problems, since a portion of the network and the computational resources will be dedicated to the security checks as users connect. Normally, solutions offering this feature use a cryptographic chip — the trusted platform module (TPM) [11] — that may pose a performance bottleneck while issuing the required verifications. SECURED overcomes this issue by introducing a trusted third party (TTP) system (Fig. 3). This is an entity that is trusted by users and infrastructure administrators, which asynchronously attests a set of controlled NEDs in a configurable time interval. The advantage of this approach is twofold. First, the workload for the attestation process does not increase with the number of connecting users, since the NED is common to all users. Second, end users will get a response regarding the integrity of the NED almost immediately.

We have developed a prototype that uses *strongSwan* [12] for the creation of a trusted channel with IPsec. To this end, strongSwan has been adapted to generate RA requests to the TTP, and either continue or drop the connection depending on the result of the integrity verification. The TTP has been implemented with *OpenAttestation* [13], a framework for attesting large infrastructures. Our initial results show that the establishment of an IPsec connection without attestation is very fast (around 76 ms), and the asynchronous attestation with the TTP in the same setting does not introduce noticeable delays (around 217 ms). Unlike our solution, performing synchronous RA adds a significant delay on the creation time of the tunnel (around 4.119 s).

Another source of overhead is due to the size of the integrity reports. Figure 5 shows the size of the reports exchanged between the NED and the TTP. The results were obtained over a 10-min period, where a user repeatedly connected to the NED. While the first report generated is near 300 kB, subsequent reports are very small (between 4 and 8 kB) due to the fact that *OpenAttestation* only sends new integrity measurements, which are performed on the NED with the Integrity Measurement Architecture (IMA) [14] software. Note that the first report contains all the measurements performed at boot time. Furthermore, new reports will be generated only if new measurements are produced on the NED (i.e., when new software is executed).

These initial results show that smartly performing RA does not incur a noticeable overhead for the end user, as all the heavy lifting is asynchronously performed behind the scenes. The analysis also sheds light on the feasibility of enabling end users to remotely verify the status of a NED. It is worth highlighting that the interval between two consecutive attestations can be configured, thereby offering the possibility of defining convenient trade-offs depending on the case. So far, we have seen that the RA of a single NED will introduce negligible overhead.

However, performing the RA over a distributed infrastructure poses complex challenges and remains an open problem. These challenges increase when we also include multi-domain scenarios or requirements such as user mobility and roaming. Furthermore, the assessment of time bounds for dynamic service deployment, as well as the appraisal of the multi-tenant isolation model, will need to be deeply analyzed in the near future. We plan to develop a comprehensive prototype that will address these issues. Our research and future evaluations will prioritize the following aspects: security and isolation, ease of use, deployment and service provisioning in relatively short timescales, and, related to the latter, support for user mobility.

User-Centric Policy Framework

Our policy-based framework also needs an in-depth performance assessment to evaluate if the policy services can actually be used in real sce-

These initial results show that smartly performing RA does not incur a noticeable overhead for the end user, as all the heavy lifting is asynchronously performed behind the scenes. The analysis also sheds light on the feasibility of enabling end users to remotely verify the status of a NED.

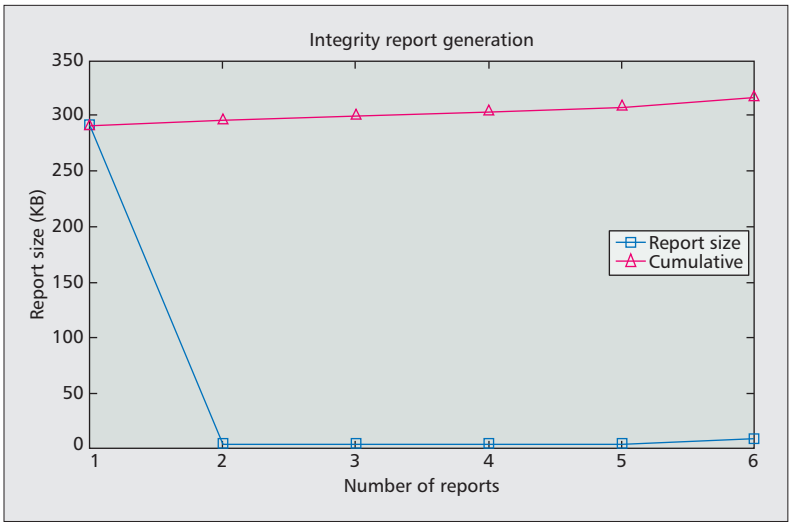

Figure 5. Size of the integrity reports generated with OpenAttestation.

narios. To this purpose, we tested the performance of the reconciliation, anomaly analysis, and translation with an off-the-shelf computer equipped with an Intel processor i7-3630QM (2.4 GHz), with 16 GB of RAM, running OpenJDK RE 1.7.0 55 on top of a Linux operating system. We performed two different rule processing experiments: an average case with a realistic amount of rules, and a higher bound worst case scenario with thousands of rules. In both cases, we considered two types of filtering within the PSAs: a *packet filter* and an *L7 filter*. During the experiments, we measured the time required to process and validate the filtering rules. As discussed earlier, such validation is composed of three parts: anomaly analysis, reconciliation, and M2L translation.

The first tests evaluate the performance of a small/medium scenario, where the number of rules per user are on average in the range of tens or hundreds. This estimation was derived from a use case with four actors, where policies included 10 to 50 rules for each PSA, amounting to an average of 100 rules to be processed. We consider that these numbers per user are representative of a reasonable average, since in a user-centric approach, the size of the rule set will not raise to thousands — which is typically the case found on border firewalls of large companies. As reported in the first row of Table 1, all three measured policy-related tasks were completed in less than 1 ms.

The second experiment aims to assess the scalability for large-scale policy scenarios. This means scenarios that, as stated on [8], statistically satisfy significant parameters of the policies that can be found in practice. This experiment provides two different results. On one hand, we compute the necessary processing time for a very large amount of rules. On the other hand, we compute the amount of rules that can be processed in 1 s — a amount for interactive purposes. Both results are reported in Table 1. We observe that for the anomaly analysis, our prototype can process 5000 rules in 12 s for the packet filtering case. In contrast, L7 filtering requires 90 s to perform the same task, due to the massive usage of regular expressions. In terms of the number of rules processed in less than a second, we obtained 2000 rules for the packet filter case and 1000 for L7 filtering. Regarding the reconciliation part, we were able to process 1500 packet filter policies and 1000 L7 filter policies in less than 1 s. However, the worst cases for the 5000 rules considered yielded reconciliation times of 74 s and 364 s for the packet filter and L7 filter, respectively. Finally, the translation of MSPL into low-level configurations is a linear problem that took approximately 1 s with 5000 rules with both an XSLT-based approach and a SAX-based Java program. All these results are summarized in Table 1.

Given that these computations are performed at infrastructure elements, wherein computational power can be adjusted as needed, we consider that our approach can reasonably scale in several real scenarios. For instance, the average cases are representative of residential scenarios, and all computations can be resolved online. We also consider that the processing of 5000 rules is quite representative of a corporate user case (e.g., an SME), and the worst cases are highly unlikely to occur in practice. Anyway, the bounds found indicate that there are cases in which the reconciliation cannot be handled online, and therefore, this analysis serves as a starting point for investigating new strategies and optimizations.

Conclusion

In this article, we have argued that for the large majority of Internet users, the current protection model against security threats is broken. Users typically have multiple devices, but achieving the same level of protection irrespective of the device used has become "mission impossible." We have proposed a paradigm shift in user protection through a user-centric model that also decouples security from user terminals. The protection model we envision is based on the setup of a trusted virtual domain per user, placed in the access network. Our approach facilitates security policy configuration, and enables uniform protection independent of the terminal used. We have also shown that the trust and security verification mechanisms offered by a prototype implementation can be applied in many practical scenarios, such as the case of residential users.

In spite of this, several of the issues addressed in this article require significant efforts in terms of research. The list is large, and includes aspects such as remotely attesting distributed systems, multi-domain scenarios (i.e., the interplay among different ISPs), user mobility and roaming scenarios, scalability analysis, assessment of upper bounds for dynamic service deployment, isolation assessment, development of a comprehensive threat model, constraints and deeper analysis of corporate scenarios, and more.

Acknowledgment

The research described in this article is part of the SECURED project [5], co-funded by the European Commission under the ICT theme of FP7 (grant agreement no. 611458).

	Filtering level	Anomaly analysis	Reconciliation	M2L translation
Average case (time to process 100 rules)	Packet filter	< 1 ms	< 1 ms	< 1 ms
	L7 filter	< 1 ms	< 1 ms	< 1 ms
Worst case (time to process 5000 rules)	Packet filter	12 s	74 s	< 1 s
	L7 filter	90 s	364 s	< 1 s
Number of rules processed in 1 s	Packet filter	2000	1500	> 5000
	L7 filter	1000	1000	> 5000

Table 1. Results of the tests for policy-based tasks.

The protection model that we envision is based on the setup of a trusted virtual domain per-user, placed in the access network. Our approach facilitates security policy configuration, and enables uniform protection independently of the terminal used.

References

[1] ETSI, "Network Functions Virtualisation (NFV) Architectural Framework," http://www.etsi.org/deliver/etsi_gs/NFV/001_099/002/01.02.01_60/gs_NFV002v010201p.pdf, December 2014.

[2] J. Sherry *et al.*, "Making Middleboxes Someone Else's Problem: Network Processing as a Cloud Service," *Proc. ACM SIGCOMM 2012 Conf. Applications, Technologies, Architectures, and Protocols for Computer .*, 2012, pp. 13–24.

[3] K. Goldman, R. Perez, and R. Sailer, "Linking Remote Attestation to Secure Tunnel Endpoints," *Proc. 1st ACM Wksp. on Scalable Trusted Computing*, 2006, pp. 21–24.

[4] OASIS, "eXtensible Access Control Markup Language (XACML) Version 3.0," http://docs.oasis-open.org/xacml/3.0/xacml-3.0-core-spec-os-en.pdf, Jan. 2013.

[5] "Security at the Network Edge (SECURED)," http://www.secured-fp7.eu/.

[6] C. Basile *et al.*, "Ontology-Based Security Policy Translation," *J. Info. Assurance and Security*, vol. 5, no. 1, 2010, pp. 437–45.

[7] IBM Global Technology Services, "IBM Security Services 2014 Cyber Security Intelligence Index," http://media.scmagazine.com/documents/82/ibm_cyber_security_intelligenc_20450.pdf, June 2014.

[8] C. Basile, A. Cappadonia, and A. Lioy, "Network-Level Access Control Policy Analysis and Transformation," *IEEE/ACM Trans. Networking*, vol. 20, no. 4, 2012, pp. 985–98.

[9] P. McDaniel and A. Prakash, "Methods and Limitations of Security Policy Reconciliation," *ACM Trans. Info. System Security*, vol. 9, no. 3, Aug. 2006, pp. 259–91.

[10] C. Basile *et al.*, "A Formal Model of Policy Reconciliation," *Proc. 23th Euromicro Int'l. Conf. Parallel, Distributed, and Network-Based Processing*, March 4–6, 2015.

[11] H. Zhang, Z. Qin, and Q. Yang, "Design and Implementation of the TPM chip J3210," *Proc. 2008 Third Asia-Pacific Trusted Infrastructure Technologies Conf.*, 2008, pp. 72–78.

[12] A. Steffen, "strongSwan — IPsec for Linux," https://www.strongswan.org.

[13] Intel, "OpenAttestation SDK: A SDK for Remote Attestation," https://github.com/OpenAttestation/OpenAttestation.

[14] R. Sailer *et al.*, "Design and Implementation of a TCG-based Integrity Measurement Architecture," *Proc. 13th Conf. USENIX Security Symp.*, 2004, pp. 223–38.

Biographies

Diego Montero (dmontero@ac.upc.edu) received his B.Sc. in computer engineering from the University of Cuenca, Ecuador. He completed his M.Sc. in computer architecture, networks and systems from the Technical University of Catalonia (UPC), Spain. He is currently a Ph.D. candidate at the Networking and Information Technology Lab (NetITLab), where his research interests include network security, SDN, network virtualization, and mobility.

Marcelo Yannuzzi (mayannuz@cisco.com) received a degree in electrical engineering from the University of the Republic, Uruguay, and MSc. and Ph.D. degrees in computer science from the Department of Computer Architecture, UPC, Spain. He is with the Corporate Technology Group at Cisco Systems International, Switzerland. Before that, he was head of NetITLab, as well as the Advanced Network Architectures (ANA) research group at UPC. He has led several projects in close interaction with European and U.S. companies and research centers. His interests lie in the areas of fog computing, IoE, security, NFV, orchestration and management, and mobility.

Adrian Shaw (adrian.shaw@hp.com) is a research scientist at HP Labs, Bristol, United Kingdom, where he is a member of the Embedded Control Points (ECP) group in the Security and Cloud Lab. His research covers the areas of operating systems, virtualization, and trusted computing. He received his M.Sc. in computer security from the University of Birmingham, United Kingdom.

Ludovic Jacquin (ludovic.jacquin@hp.com) is a research scientist at HP Labs in the Security and Cloud Laboratory at Bristol. He received his Ph.D. in computer science from Grenoble University in 2013, with a thesis titled *Performance/Security Trade-off for High-Bandwidth Internet VPN Gateways* supervised by Vincent Roca and Jean-Louis Roch. His main research currently focuses on infrastructure security, especially attestation of the network devices in the new "softwarized" paradigm.

Antonio Pastor (apastor@tid.es) is a technology expert in security on networks working for the network virtualization group in the GCTO unit within Telefónica I+D. Since 2006 he has been working as an expert in IP network security designs and services, and holds several certifications from ISACA and GIAC in this area. Currently, he is working on SDN and NFV technologies oriented toward security, including applied research projects and close-to-market services.

René Serral-Gracià (rserral@ac.upc.edu) received his degree in computer science (2003) and a Ph.D. (2009) from UPC. He is the R&D head of NetITLab at UPC, where he is leading different research initiatives, including projects under the European FP7 Research Framework as well as with industry. He is also an associate professor in the Department of Computer Architecture at UPC. His research interests are focused on SDNs, overlay networks, network security, routing optimization, and QoE assessment of multimedia traffic.

Antonio Lioy (lioy@polito.it) (M.Sc. in electronic engineering and Ph.D. in computer engineering) is a full professor at the Politecnico di Torino, Italy, where he leads the TORSEC group. His current research interests are network security (especially optimization and automatic configuration), PKI applications (e-identity and digital workflows), and policy-based protection of ICT systems. He is the coordinator of the SECURED project, and frequently acts as a cybersecurity expert for the Italian government and the European Commission.

Fulvio Risso (fulvio.risso@polito.it) received his Ph.D. degree in computer and system engineering from Politecnico di Torino in 2000. He is currently an assistant professor with the Department of Control and Computer Engineering, Politecnico di Torino. His current research activities focus on efficient packet processing, traffic analysis, and programmable networks.

Cataldo Basile (cataldo.basile@polito.it) received his Ph.D. degree in computer and system engineering from Politecnico di Torino in 2005. He is currently a research assistant at Politec-

nico di Torino. His research is concerned with policy-based management of security in networked environments, policy refinement, general models for detection, resolution and reconciliation of specification conflicts, and software security.

ROBERTO SASSU (rsassu@suse.de) is a senior engineer at SUSE Linux GmbH. Previously, he worked as a research assistant at Politecnico di Torino. His research activity focused on sensitive data protection, platform integrity evaluation, and cloud computing. He received his M.Sc. in computer engineering from Politecnico di Torino.

MARIO NEMIROVSKY (mario.nemirovsky@bsc.es) received a Ph.D. degree in electrical and computer engineering from the University of California, Santa Barbara in 1990. He was an adjunct professor at the same university from 1991 to 1998. After being chief architect at companies such as Apple, Inc., National Semiconductors, and General Motors (GM), he founded several renowned startups including XStream Logic, FlowStorm Networks, ConSentry Networks, and Miraveo. In 2007, he became an ICREA Senior Research Professor with the Barcelona Supercomputing Center (BSC), Spain. He holds more than 60 issued patents. His current research interests include multithreaded multicore systems, high-performance systems, IoT, big data, and network processors.

FRANCESCO CIACCIA (francesco.ciaccia@bsc.es) is currently a researcher at BSC. He is part of the Unconventional Computer Architecure and Networks research group. He received his M.S. degree in computer engineering from Politecnico di Torino. His main reaserch interests include network security, SDNs, virtualization, and the Internet of Things.

MICHAEL GEORGIADES (michaelg@prime-tel.com) is the R&D manager at PrimeTel PLC and an adjunct faculty member at the Open University of Cyprus. He received a B.Eng. from King's College London (2000), an M.Sc. from University College London (2001). and a Ph.D. from the University of Surrey (2008) in telecommunications. He has been involved in more than 10 EU ICT projects, and published more than 40 journals, book chapters, and conference publications. He received the Nokia Prize of Research Excellence for a Patent in 2004.

SAVVAS CHARALAMBIDES (savvasch@prime-tel.com) is an R&D research engineer at PrimeTel PLC. He received his B.Sc. in computer science and engineering from the University of Patras, Greece, in 2012 and his M.Phil. in Advanced Computer Science from the Computer Laboratory of the University of Cambridge, United Kingdom, in 2013. He has five publications in international journals and conferences, and is a co-author of a book chapter in the area of computer network simulations.

JARKKO KUUSIJÄRVI (jarkko.kuusijarvi@vtt.fi) received B.Sc. (Tech.) and M.Sc. (Tech.) degrees from the University of Oulu, Finland, in 2008 and 2010, respectively. Since 2010, he has been a research scientist at VTT. His current areas of research interests include mobile applications, security visualization, and cybersecurity.

FRANCESCA BOSCO (bosco@unicri.it) is a UN project officer working in UNICRI responsible for developing research and capacity building activities on misuse of technology and technology-enabled crimes. She is a member of the Advisory Groups on Gender and Secure Societies in the framework of Horizon2020 and of the Internet Security Expert Group of the EC3. She is a co-founder of the Tech and Law Center and a member of the Centre for Internet & Human Rights of European University Viadrina.