

Secure Routing for Future Communication Networks

Edited by

Amir Herzberg¹, Matthias Hollick², and Adrian Perrig³

1 Bar-Ilan University – Ramat Gan, IL, amir.herzberg@gmail.com

2 TU Darmstadt, DE, mhollick@seemoo.tu-darmstadt.de

3 ETH Zürich, CH, adrian.perrig@inf.ethz.ch

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 15102 “Secure Routing for Future Communication Networks”. Routing is a fundamental mechanism in communication networks, and its security is critical to ensure availability and prevent attacks; however, developing and deploying secure routing mechanism is still a challenge. Significant research effort is required to advance routing security in key areas: intra-domain routing, inter-domain routing, routing in new Internet architectures, and routing in mobile and wireless networks. The seminar covered these general aspects along with the following important guiding questions. How to systematise the topic area of routing security? What are evolutionary or revolutionary options towards more secure routing systems? How to secure inter-domain routing? How to secure intra-domain routing and routing in mobile/wireless settings? How to achieve data plane/forwarding security?

Seminar March 1–4, 2015 – <http://www.dagstuhl.de/15102>

1998 ACM Subject Classification C.2.0 General, C.2.1 Network Architecture and Design, C.2.2 Network Protocols

Keywords and phrases Security, Secure routing, Communication networks, Future internet, Mobile and wireless networks

Digital Object Identifier 10.4230/DagRep.5.3.28

Edited in cooperation with Michael Noisternig

1 Executive Summary

Amir Herzberg

Matthias Hollick

Adrian Perrig

License  Creative Commons BY 3.0 Unported license
© Amir Herzberg, Matthias Hollick, and Adrian Perrig

Routing is a fundamental mechanism in communication networks, and its security is critical to ensure availability and to prevent attacks; however, developing and deploying secure routing mechanisms is still a challenge. Routing is the process by which information is passed via the communication network, from source to destination, via a series of intermediary nodes/routers. Routing attacks include route-hijacking, i.e., diverting traffic to an adversary-controlled router, and denial-of-service attacks exploiting the routing mechanism, i.e., preventing communication (in parts or the entire network), e.g., by malicious dropping of packets by a router.

Routing, and even more secure routing, are complex problems with many variants. In particular, the Internet is a federation of many domains (usually referred to as autonomous systems (ASes)), each managed by a separate organization; there are separate standard



Except where otherwise noted, content of this report is licensed under a Creative Commons BY 3.0 Unported license

Secure Routing for Future Communication Networks, *Dagstuhl Reports*, Vol. 5, Issue 3, pp. 28–40

Editors: Amir Herzberg, Matthias Hollick, and Adrian Perrig



DAGSTUHL
REPORTS Dagstuhl Reports

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

protocols for routing inside an AS (intra-domain routing) and for routing from a source in one AS to a destination in a different AS (inter-domain routing). Significant efforts are dedicated to securing intra-domain routing protocols and inter-domain routing protocols; in addition, significant efforts are also dedicated to the design of completely new Internet architectures that include secure routing mechanisms.

Another categorization of routing mechanisms and challenges involves mobility. Many routing protocols, including standard Internet routing, are designed for a mostly static topology, where connections between routers are relatively stable. However, communication is increasingly performed among mobile devices. There are many efforts and challenges in the design of (secure) routing mechanisms for highly mobile networks, e.g., between tiny wireless sensors, swarms of tiny robots, or simply mobile users (e.g., upon catastrophic failure to regular infrastructure).

There is also a need to re-evaluate and possibly re-design routing mechanisms and security measures, to address changes in the way the Internet is used, and in the presence of new security challenges. In particular, is there a need to adapt routing to facilitate, and/or take advantage of, cloud services, and to support security for them? Is there a need to adapt routing to the increased threat of Denial-of-Service attacks, or to facilitate widespread provision of Quality-of-Service? Should routing be modified to take into account energy considerations, or to take advantage of and facilitate Software Defined Networking (SDN)? If modifications are made for these goals, how does this affect routing systems' attack surface? Finally, is there a need to modify routing and its security mechanisms, as a result of the recent revelations regarding the scope of abuse of routing by powerful nation-state adversaries?

In summary, to advance routing security in the aforementioned topic areas, a number of significant research problems need to be addressed, and identifying these problems was the goal of this seminar. The first objective was to facilitate brainstorming and exchange of ideas among experts working in different areas and types of secure networking, leading to an improved understanding of the different aspects of secure routing. The second objective was to identify the most important research challenges and to devise a roadmap towards addressing urgent issues. Through the seminar, we aimed at opening up new avenues of research in the area of routing security. For the given focus areas of the seminar, we contributed to the following key research challenges:

- Routing Security by Design for a Future Internet: the challenge was to overcome the limitations and confined models imposed by today's Internet. Both clean slate as well as evolutionary approaches towards a secure-by-design future Internet were discussed.
- Inter-domain Routing Security and Intra-domain Routing Security: challenges addressed in inter-domain routing were the reconciliation of potentially conflicting security interests across multiple domains and resilience against recently published attacks. Intra-domain routing is underrepresented in research; here, the seminar aimed at identifying the key research challenges towards a research roadmap.
- Routing Security in Mobile/Wireless Networks, and in Delay- and Disruption-tolerant Networks: the main goal within the seminar was to identify possible ways to provide routing security in light of the severely limited resources and special characteristics of mobile and wireless systems.
- Quality of Service (QoS) and Denial of Service (DoS) aspects of Routing Security: the challenge was to jointly consider security considerations and QoS aspects, both in theory and practice.

To address these challenges, the seminar was organized in six working groups. They are presented in Section 4 of this report. The schedule of the seminar and its working groups is presented in Table 1 below.

■ **Table 1** Schedule for Dagstuhl Seminar 15102.

	Sunday	Monday	Tuesday	Wednesday
Breakfast from 7:30–8:45				
9:00–10:30		+ Welcome by Matthias + Intro of participants + Intro talk by Steven	Wrap-up & update Plenary discussion: A, B, F	Wrap-up & update Plenary discussion: C, D, E
Coffee				
10:50–12:10		+ Talk by Adrian Perrig + Talk by Randy + Agenda, plan, goals by Amir Picture of group	Parallel sessions C – Inter-domain D – Mobile/Wireless E – QoS/DoS, Forwarding security	Wrap-up & reserved
Lunch (12:15)				or lunch boxes for early departure
13:30–15:30		Parallel sessions A – Taxonomy B – (R)Evolution	Social event: 1h hiking/biking around Dagstuhl castle from 14:30 Parallel sessions contd.: C – Inter-domain D – Mobile/Wireless E – QoS/DoS, Forwarding security	Departure
Coffee	Arrival, registration from 15:00 to 19:00			
15:50–17:50		Parallel sessions contd.: A – Taxonomy F – Intra-domain	Parallel sessions contd.: C – Inter-domain D – Mobile/Wireless E – QoS/DoS, Forwarding security	
Dinner (18:00)				
	Cafe, wine, cheese, discussions in cafeteria, wine cellar, music room, etc.			

WG (A) Towards a taxonomy on secure routing

WG (B) Revolution and/or evolution?

WG (C) Securing inter-domain routing

WG (D) Routing security in mobile/wireless networks and delay-/disruption-tolerant networks

WG (E) Forwarding/data-plane security

WG (F) Intra-domain routing security

2 Table of Contents

Executive Summary

Amir Herzberg, Matthias Hollick, and Adrian Perrig 28

Overview of Talks

Routing is as Insecure as the Rest of the Flippin' Internet, but it's Scarier
Randy Bush 32

Routing Security Challenges
Steven Bellovin 32

SCION: A Secure Next-Generation Internet Architecture
Adrian Perrig 32

Working Groups

Towards a Taxonomy on Secure Routing 33

Revolution and/or Evolution? 35

Securing Inter-Domain Routing 36

Routing Security in Mobile/Wireless Networks and Delay-/Disruption-Tolerant
Networks 37

Forwarding/Data Plane Security 37


Intra-Domain Routing Security 38

Participants 40

3 Overview of Talks

3.1 Routing is as Insecure as the Rest of the Flippin' Internet, but it's Scarier


Randy Bush – Internet Initiative Japan Inc., JP

License  Creative Commons BY 3.0 Unported license
© Randy Bush

The goal of the opening talk is to raise lessons from other (Internet) evolutions and revolutions to make us aware of the pitfalls such as 'second-system syndrome incompatibility'. Routing protocols need to protect their assets (traffic content and meta data) from threats such as traffic content inspection, modification, injection, and analysis. Both routing and DNS attacks are today's most severe security issues on the Internet, but routing threats are on the rise. Attackers are primarily spammers, but also governments and financial institutions are involved. Routing attacks target external infrastructure (e.g., IRR, Whois, RPKI) and both well-implemented and poorly designed protocols. This is a disaster happening every day but a cure is difficult to deploy. Security solutions need to provide a real benefit to some involved entities, and deployment must be simple and backwards compatible. Lessons may be learnt from bad examples such as IPv6 and better ones such as RPKI.

3.2 Routing Security Challenges

Steven Bellovin – Columbia University, US

License  Creative Commons BY 3.0 Unported license
© Steven Bellovin

Routing security is hard because failures occur when someone lies. The protocol, however, is executed correctly, and the liar may be distant. Attackers today are definitely spammers, but governments may be involved. Attackers can lie about prefixes for paths. It can happen internally or externally. There are currently no good incentives for deployment in the interdomain case.

3.3 SCION: A Secure Next-Generation Internet Architecture

Adrian Perrig – ETH Zürich, CH

License  Creative Commons BY 3.0 Unported license
© Adrian Perrig

The Internet has been successful beyond even the most optimistic expectations. This success has created a dependency on communication. Unfortunately, the current Internet suffers from numerous vulnerabilities and shortcomings that limits its availability. To address these issues, we study the design of a next-generation Internet architecture that is secure, available, and offers privacy by design; that provides incentives for a transition to the new architecture; and that considers economic and policy issues at the design stage.

4 Working Groups

Core topics of the seminar have been organized in six working groups. Each working group lasted for one or two sessions of 120 minutes each and was running with one or two other working groups in parallel (one working group was merged with another one due to a low number of participants). Working group sessions were followed by a wrap-up session the next day, in which the outcome was presented to all seminar participants. The following subsections provide a summary of the discussions within each working group.

4.1 Towards a Taxonomy on Secure Routing

This working group has explored how secure routing protocols could be categorized towards a better understanding of existing solutions and the areas that need improvement. Particular focus has been put on identifying goals of secure routing protocols as this is an area that currently is poorly specified.

Principal classifiers for secure routing protocols that have been identified include

- attacks/vulnerabilities/risks on the routing protocol,
- security solutions,
- the routing abstraction/model/formalization/environment used, and
- goals at the control and data plane.

It is important that goals are clearly specified as part of a system specification. In terms of routing abstractions/models/formalizations, one may consider

- information-centric/content-centric networking (e.g., in Future Internet),
- software-defined networking (makes secure routing even more complex, but at the same time there is a chance to do things exactly as they should be),
- traditional Internet routing (intra- and inter-), and
- mobile/wireless routing.

These properties form what was called the environment of the routing protocol.

Goals

The participants have determined that a taxonomy of current routing protocols based on goals is difficult because today there are no formalisms for (routing) protocols. People talk about safety, aliveness, lightness, etc., but these issues are not defined, and models do not specify what is expected from networks. For example, a goal such as availability seems clear, but under which circumstances is usually not specified. Routing protocols typically provide a best-effort service under non-adversarial circumstances but specifications are not very clear what can be expected under which (threat) scenarios. This issue has been identified by the working group as an open problem on the network specification/modelling side.

Thus, what is needed is a model that clearly specifies what to be expected. To this end, goals need to be separately specified for the data, control, and management plane, and they may be different for different network entities (such as source vs. intermediate node vs. destination) defined by the model.

The following goals or requirements for a secure routing protocol have been identified, distinguishing the data plane from the control plane and the management plane:

- Data plane:
 - It should provide a certain service, which may be best-effort, reliable communication, or reliable communication with performance guarantees.
 - It may provide (end-to-end) path enforcement.
 - There are a number of common security/protection goals, which include confidentiality, integrity, and authenticity (CIA) as well as privacy.
 - Other goals that have been identified included net neutrality and censorship resilience. What is meant by the latter needs to be better understood.
- Control plane:
 - We may distinguish between two aspects:
 - * centralized (e.g., SDN) vs. decentralized approaches, and
 - * whether the routing protocol implements local state, global state, or is stateless. Examples include distance-vector routing, which is decentralized with local routing state, and link-state routing, which is decentralized but acquires a global state.
 - Common security goals include confidentiality, integrity, and authenticity (which may be different between the control and the data plane).
 - Freshness and trustworthiness/completeness for a global routing state. Trustworthiness in this context is not about integrity but about making sure that we can trust the global state that we aggregate.
- Management plane:
 - It must be reliable and timely.
 - In general, goals are expected to be similar to those of the data plane.

The control plane may cover one or more administrative domains, which needs to be taken into account. At the same time, there are routing protocols without a control plane (e.g., flooding). The management plane, which is used mostly for monitoring the data and control planes (but also for configuration and reboot), may implement message routing via dedicated wires or some ‘obfuscated’ mechanism and may be multi- or single-hop. Attacks on the management plane are similar to those on the data plane though the constraints are tighter.

A worthwhile reference that was mentioned is the paper [1], which provides a taxonomy of common concepts and definitions for security goals in communication. The paper also motivates looking at the boundary between the routing system and the environment. For example, the attack surface depends on system boundaries, e.g., malicious hardware of router, or attacks at interfaces such as power-attack/connection to 3G/4G. Another important issue at the boundary is how to bootstrap the crypto (i.e., how key management is provided).

In terms of threats, it may be interesting to look at a study published by the European Union Agency for Network and Information Security (ENISA) [2], which summarizes good practices that aim at securing an Internet infrastructure asset from Important Specific Threats. The study includes a mind map, which graphically categorizes Internet infrastructure assets into eight families: protocols, software, hardware, information, human resources, facilities, interconnection, and services; the latter which is split into four subfamilies: applications, routing, addressing, and security.

A side question was whether/how vendors would be incentivized to work with the community to implement these goals. While this has been identified as a separate issue, a clear specification of goals would at least allow checking whether the protocol actually provides these goals.

Solutions

Participants in the working group agreed that it is not possible to provide a concise and exhaustive list of solutions that implement the goals that have been identified. Therefore, only a few examples have been discussed. It was assumed that any mechanisms for the data plane must be able to rely on the control plane ‘doing its job’. Examples discussed included:

- Confidentiality, integrity, and authenticity can be provided by classical cryptography, both at the control and data plane. Attacks are the ‘usual suspects’ (side-channel attacks, etc.).
- Reliability with performance guarantees needs to be implemented by
 - Reliability measures such as forward error correction, network coding, and (negative) acknowledgments.
 - Performance measures, which include prioritization, admission control, and resource reservation.
- Trustworthiness/completeness needs to be based on cryptography, reputation mechanisms, heuristics, and invariants.

Future Internet

Goals for the Future Internet have only been briefly discussed. Comments included the need to consider non-interference for parallel stacks and to be aware of downgrade attacks when backwards compatibility is provided, and that routing over hybrid networks may be a challenge.

Outcome

It was decided that it would be a good idea to write a survey paper on this topic. Currently, there are surveys on routing taxonomy, but they are solely on wireless routing. Some older survey that focuses on BGP exists but it does not cover newer developments such as software-defined networking. Reference [3] is also not recent and it does not seem very systematic.

In terms of the content of the paper, it was suggested that it needs to cover (1) the system environment (akin to a trusted computing base) and its boundaries, (2) goals in terms of functionality and security on data, control, and management plane, which then allows us to describe (3) mechanisms to fulfil the goals, and subsequently (4) possible attacks on the mechanisms. It was also suggested that we should consider the industry side, which sees several security issues as rather academic from the viewpoint that detection is already good enough while protection at the cost of overhead is not worthy. From this perspective, it was agreed that the paper needs to include a clear and realistic adversary model.

4.2 Revolution and/or Evolution?

This working group was concerned with the problem of incrementally deployable improvements in secure communication and routing. It has raised two important issues towards addressing the problem:

- The identification of long-term ideal situations (the “vision”). This bears the question whether the vision is realistic.
- The identification of intermediate steps. The questions here are: what are incentives; and what is the cost?

There are a number of threats in incremental deployments. These include: (1) Not having a vision. The vision may be too vague, too tight, or too ambitious. (2) Not all intermediate steps constitute an improvement ([4]), and partial deployment might even lead to vulnerable conditions. (3) Partial deployment may be hard to achieve, e.g., ingress filtering only prevent others, so there may be no real incentives. (4) Corporate interests may result in conflicting goals.

Ultimately, the goal is to implement the vision, which includes identifying and implementing all intermediate steps.

Studying current deployments, the following systems have been identified: LISP, PKI, spoofing prevention, secure E2E communication (SSH, IPsec, SSL), anonymous communication (Crowds, TOR), digital currencies (cybercash, café, Mondex, Millicent, Bitcoin), secure email, and public-key crypto.

In terms of research activities, the Internet architecture board (IAB) recently held a workshop on Internet technology adoption and transition and published its findings in RFC 7305 [5].

The working group concluded with identifying a challenge for future work, which is to find other incentives for incremental deployments than just the operator's economic ones.

4.3 Securing Inter-Domain Routing

Secure route discovery for inter-domain routing faces a number of challenges, including:

- Privacy of control-plane data (e.g., of relationships between operators)
- Correctness of the control plane (e.g., hijacking of routes must be prevented, i.e., routers must know accurate, trustworthy routes to (all) destinations; and ownership must be asserted)
- Deployment (costs and incentives need to be considered; this may be even more challenging for partial deployment)
- Convergence (stable state vs. optimal state, efficiency)
- Policy (source / destination / transit)
- Trust relationships
- Validation of identifiers

These challenges are not the least due to pressing security threats. Examples include: (1) hijacking of resources (identifiers, links, IPs), (2) topology attacks (path redirection, link cutting), (3) availability attacks (protocol, DDoS, complexity attacks), and (4) privacy attacks (“learning” attack).

To address the challenges, a number of goals have been identified:

- Functionality: this is about reaching equilibrium for convergence.
 - A specific question concerns how to characterize such equilibria (correctness), e.g., correct path establishment to ensure correct packet delivery (if a valid path exists, it should be found), loop-freeness, and prevention of invalid name announcements.
 - Convergence must even be achievable under attack.
- Trust: creation (who do you trust for what), representation, distribution, management, agility & flexibility.
- Policies for source and destination, and for transit.
- Security in case of partial deployment.
- Privacy
 - of the topology, and
 - of the policy.

The working group has then discussed the current state-of-art in secure inter-domain routing in terms of deployment and research. The following deployed systems have been identified: RPKI, Origin Authentication (coming up), and BGPsec (close to being finished)

For BGPsec, a number of problems have been highlighted: RPKI and certificate managements, hierarchies, resource intensive (crypto, memory), DoS (fake withdrawal attacks), partial deployment, convergence, and deployment incentives.

In terms of research, several activities within the context of Future Internet research efforts can be found: (1) SCION (which provides incentives for national/local deployment), (2) FIA, (3) MobilityFirst, (4) ILNP, (5) LISP, and (6) HIP.

The working group concluded with an identification of the following research challenges:

- interdomain multipath routing (beneficial for security)
- diverse paths
- non-hierarchical trust (toleration of malicious nodes)
- impact of security on the system (measurements, verification)
- guaranteed policy-compliant path computation (if path exists, it should be found, understood, and useable)
- centralization of route computation
- finding policy-compliant path
- building network model out of specifications and the formal statements about network properties
- deployment incentives
- routing transparency
- control plane congruence with data plane
- proof of ownership in hierarchy

4.4 Routing Security in Mobile/Wireless Networks and Delay-/Disruption-Tolerant Networks

Due to a limited number of participants, it was decided to join a larger discussion within the working group on a taxonomy of secure routing protocols.

4.5 Forwarding/Data Plane Security

This working group has addressed the issue of providing a forwarding service under QoS requirements and DoS attacks, which has real-world application in the intra-routing domain. Aspects that need to be considered for this service include

- latency and latency variance – this needs to consider limited throughput but does not assume any adversary,
- (local) fast recovery,
- global monitoring, and
- path enforcement (under non-adversarial scenarios).

Fast recovery and global monitoring provide a reliability service but not end-to-end reliability. Current systems are decentralized but are moving towards central control (i.e., SDNs).

From a practical point of view, IPsec tunnels may be used to provide forwarding security. However, they are not used for various reasons. For example, mobile base stations can maintain just a few security associations, so this is a point of tension. Privacy is not

considered a strict requirement since a cellular operator can disable encryption because of load (e.g., consider New Year’s text message storm). However, privacy in residential applications might be a desirable marketing option.

Concerning DoS attacks and traffic analysis, current systems provide no protection against traffic analysis. Existing techniques to mitigate traffic-based DoS attacks include filtering, in particular the redirection of traffic through optimized filters. However, some defenses have undesirable effects on intra-domain routing such as when path security is needed.

A threat model for forwarding security must consider both outsiders and insiders. Attacker capabilities include:

- Outsiders:
 - Traffic injection (possibly coordinated)
 - Eavesdropping (wireless links, only)
- Insiders:
 - Packet delay, dropping, misdirection, injection, corruption (payload and header), replays, etc.
 - Attacks may be directly carried out on the data plane, whereas all attacks on the control plane impact the data plane.
 - The ENISA report on BGP [6] provides a good overview of threats.

In addition to attacks on the communication links, an adversary may attack the operating system and the hardware. These system attacks ‘on the box’ impact the data plane as an adversary may delay data, drop data, etc. (same as before). Adversaries may also collude (e.g., by setting up a wormhole between two communication devices, which amounts to a control-plane attack affecting the data plane).

Solutions for these threats must provide clear incentives (e.g., financial incentives). This is possible when a security solution is an ‘enabler’ (e.g., for QoS).

It has been mentioned that while forwarding security is foremost considered for traditional, static networks, mobile wireless networks must be addressed, too. This includes cellular networks, wireless multi-hop networks, and specific settings such as medical networks.

4.6 Intra-Domain Routing Security

Secure intra-domain routing protocols are currently underrepresented in both the academic literature and in deployed systems. The reasons for this unfortunate fact are not obvious; attempts to identify why neither researchers nor industry have sufficiently addressed this important problem resemble a guessing game. In existing intra-domain protocols, however, attackers can easily create damage by launching various attacks.

The problem statement in this domain is thus how to secure intra-domain routing and make it robust against malicious entities, be it routers, end-hosts, or administrators. Such entities may launch a number of attacks including:

- Redirection attacks
 - Path shortening (LSA alteration, bogus links)
 - Prefix announcements
 - Link cutting or congestion
- Availability attacks (delay LSA or drop)

A practical problem that has been identified is that incidents are often not reported (in comparison to the inter-domain case). A first step in addressing intra-domain routing attacks is the deployment of better monitoring techniques.

Acknowledgments. We thank Raphael Reischuk for acting as a scribe during the Dagstuhl seminar and Michael Noisternig for acting as a scribe during the seminar and the collector for this report.

References

- 1 Avizienis et al., *Basic Concepts and Taxonomy of Dependable and Secure Computing*. IEEE Trans. on Dependable and Secure Computing, vol. 1(1), pp. 11–33. 2004.
- 2 *Threat Landscape and Good Practice Guide for Internet Infrastructure*. European Union Agency for Network and Information Security (ENISA), Jan. 2015. <https://www.enisa.europa.eu/activities/risk-management/evolving-threat-environment/enisa-thematic-landscapes/threat-landscape-of-the-internet-infrastructure>
- 3 Chakrabarti et al., *Internet Infrastructure Security: A Taxonomy*. IEEE Network, vol. 16(6), pp. 13–21. 2002.
- 4 Lychev et al., *BGP security in partial deployment: is the juice worth the squeeze?*, ACM SIGCOMM Computer Communication Review, vol. 43(4), pp. 171–182. Oct. 2013.
- 5 Lear (ed.), *Report from the IAB Workshop on Internet Technology Adoption and Transition (ITAT)*, IETF RFC 7305. July 2014.
- 6 *Secure routing: State-of-the-art deployment and impact on network resilience*. European Network and Information Security Agency (ENISA), July 2010. <https://www.enisa.europa.eu/publications/archive/state-of-the-art-deployment-and-impact-on-network-resilience>

Participants

- Steven Bellovin
Columbia Univ. – New York, US
- Saleem Bhatti
University of St. Andrews, GB
- Randy Bush
Internet Initiative Japan Inc. –
Tokyo, JP
- Joel M. Halpern
Leesburg, US
- Amir Herzberg
Bar-Ilan University -
- Ramat Gan, IL
- Matthias Hollick
TU Darmstadt, DE
- Ivan Martinovic
University of Oxford, GB
- Rossella Mattioli
ENISA – Athens, GR
- Cristina Nita-Rotaru
Purdue University – West
Lafayette, US
- Michael Noisternig
TU Darmstadt, DE
- Panagiotis Papadimitratos
KTH Royal Institute of
Technology, SE
- Adrian Perrig
ETH Zürich, CH
- Raphael Reischuk
ETH Zürich, CH
- Alvaro Retana
CISCO Systems – Research
Triangle Park, US
- Michael Schapira
Hebrew Univ. – Jerusalem, IL
- Thomas C. Schmidt
HAW – Hamburg, DE
- Jean-Pierre Seifert
TU Berlin, DE
- Haya Shulman
TU Darmstadt, DE
- Mahesh Tripunitara
University of Waterloo, CA
- Gene Tsudik
Univ. of California – Irvine, US
- Laurent Vanbever
ETH Zürich, CH
- Matthias Wählisch
FU Berlin, DE

