

**FLAMINGO***European Seventh Framework Network of Excellence*<http://www.fp7-flamingo.eu/>

## **WP5 — Network and Service Monitoring**

***Deliverable D5.3 — Third year deliverable on network and service monitoring***

© Copyright 2014 FLAMINGO Consortium

University of Twente, The Netherlands (UT)  
Institut National de Recherche en Informatique et Automatique, France (INRIA)  
University of Zurich, Switzerland (UZH)  
Jacobs University Bremen, Germany (JUB)  
Universität der Bundeswehr München, Germany (UniBwM)  
University Politecnica de Catalunya, Spain (UPC)  
iMinds, Belgium (iMinds)  
University College London, United Kingdom (UCL)



Project funded by the European Union under the  
Information and Communication Technologies FP7 Cooperation Programme  
Grant Agreement number ICT-FP7 318488

## Document Control

**Title:** D5.3 — Third year deliverable on network and service monitoring  
**Type:** Public  
**Editor(s):** Anna Sperotto  
**E-mail:** a.sperotto@utwente.nl  
**Doc ID:** D5.3  
**Delivery Date:** 31.10.2015  
**Author(s):** Anthea Mayzaud, Abdelkader Lahmadi,  
Christos Tsiaras, Christian Dietz, Daphne Tuncer  
Marinos Charalambides, Mario Flores,  
Mario Golling, Maxim Claeys, Niels Bouten, Gaëtan Hurel  
Radhika Garg, Rashid Mijumbi, Ricardo Schmidt,  
Rick Hofstede, Sebastian Seeber, Corinna Schmitt,  
Guilherme Sperb Machado, Jair Santanna, Stefano Petrangeli, Vaibhav Bajpai

For more information, please contact:

Dr. Aiko Pras  
Design and Analysis of Communication Systems  
University of Twente  
P.O. BOX 217  
7500 AE Enschede  
The Netherlands  
Phone: +31-53-4893778  
Fax: +31-53-4894524  
E-mail: <a.pras@utwente.nl>

## Legal Notices

The information in this document is subject to change without notice.

The Members of the FLAMINGO Consortium make no warranty of any kind with regard to this document, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. The Members of the FLAMINGO Consortium shall not be held liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

## Executive Summary

Deliverable D5.3 reports on the achievements of WP5 with respect to monitoring of networks and services. In this document, we first give an overview of the activities and the research that have taken place during Y3. Then, we report a set of selected research highlights, particularly focusing on large scale measurements and measurements of security events, two of the core areas of WP5.

The S.M.A.R.T. objectives (Section B.1.1.5 of the DoW) that are key to this WP are 1) integration of Ph.D. students, and 2) producing scientific publications (see Sec. 2.1). Considering both the Ph.D. collaborations active in the project and the scientific output for Y3, we can confirm that WP5 has exceeded all S.M.A.R.T. objectives.

Besides the S.M.A.R.T. objectives, WP5 has also delivered excellent results with respect to WP5-specific objectives, which are reported in Section 2.2. Currently, research is ongoing in all the WP-specific objectives but two, which have been already targeted in previous years. Particularly active and successful is the research conducted in the context of Objective 5 ("To collect (anonymized) monitoring data"). The research conducted for fulfilling these objectives resulted in some of the top publications for WP5.

The key achievements for Y3 are:

- WP5 focused even more decisively on publishing in top conferences and journals in the fields of networking, network measurements and network and service management (Sections 4–9). This resulted in papers at, for example, IMC 2014, SIGCOMM 2015 (short paper), TNSM and CCR. The short paper at SIGCOMM 2015 describes a new measurement infrastructure that has been developed at UT in collaboration with SURFnet, aiming at collecting daily snapshots of the state of DNS by querying all domain names registered in .com, .net and .org zones, roughly equivalent to 50% of the registered domain names worldwide (Section 5). The TNSM paper investigates the impact of different packet sampling strategies on link dimensioning (Section 9). The publications in CCR are an editorial reporting on the lessons learned when using the RIPE Atlas platform for network measurements (Section 4) and a paper advocating, based on evidences supported by network measurements, the use of elliptic curve cryptography in place of RSA in DNSSEC (Section 7).

Visible evidence of the quality of those publications are, besides the venues themselves, for example the *IMC Community Contribution Award* awarded to the paper [1] for the best public dataset. The same paper has also been awarded a *IRTF Applied Networking Research Prize 2015*. In addition, several papers are currently under reviews in venues such as INFOCOM 2016, IEEE Communications Magazine and IEEE Journal on Selected Areas in Communications.

WP5 is also focusing in collaborations. For WP5 publications, 21 have been achieved in collaboration with other EU projects and institutions. In addition, some of the measurements activities, in particular the ones around DNS (the Internet of Names and the collaboration with the SAND project) have open concrete possibilities of collaborations with institutions like Center for Applied Internet Data Analysis (CAIDA, UCSD, USA) and the Information Sciences Institute (ISI, USC, USA). Long term visits have already been planned and possibilities in terms of data sharing are being investigated.

- In Y1, the project published 37 papers and in Y2, the number of publications for the overall project was 50. In Y3, we exceeded the previous year number of publications by publishing 73 papers, many of which containing a strong measurement component. In response to the reviewers comments, in Y3 the project has also mapped the publications to their most closely

related WP (see D8.3). An overview of the scientific output for WP5 is given in Section 2.1. The complete list of FLAMINGO publications can be found in D8.3.



# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Objectives and Tasks</b>	<b>2</b>
2.1	S.M.A.R.T. Objectives . . . . .	2
2.2	Workpackage Objectives . . . . .	2
2.2.1	Ongoing Objectives . . . . .	3
2.2.2	Open Objectives . . . . .	7
2.3	Tasks and Objectives Mapping . . . . .	7
<b>3</b>	<b>Integration of PhD Students</b>	<b>9</b>
3.1	PhD Student Collaborations . . . . .	10
3.2	Description of the Collaborations . . . . .	11
3.2.1	Intrusion Detection Systems (UT-UniBwM-IDS) . . . . .	11
3.2.2	Flowoid: a NetFlow/IPFIX Probe for Android-based Devices (UT-INRIA-Flowoid) . . . . .	15
3.2.3	Mobile Measurements (UZH-JUB-UniBwM-M2) . . . . .	16
3.2.4	Schengen Routing (UT-UniBwM-Scheng) . . . . .	17
3.2.5	Machine Learning and Botnet Detection (UniBwM-UT-MLB) . . . . .	18
<b>4</b>	<b>Lessons Learned from using the RIPE Atlas Platform for Measurement Research</b>	<b>20</b>
4.1	Heavy-tailed probe distribution . . . . .	20
4.1.1	Clustering probes by ASN . . . . .	21
4.1.2	Ranking ASNs by number of probes . . . . .	21
4.1.3	Clustering ASNs by network type . . . . .	21
4.1.4	Skewed distribution of probes . . . . .	23
<b>5</b>	<b>The Internet of Names: A DNS Big Dataset</b>	<b>24</b>
5.1	Infrastructure . . . . .	25
5.1.1	High-level overview . . . . .	25
5.1.2	Main measurement . . . . .	25
5.1.3	Analysis . . . . .	26
5.2	Case study . . . . .	26
5.3	Conclusions and Future Work . . . . .	26
<b>6</b>	<b>Booters – An Analysis of DDoS-as-a-Service Attacks</b>	<b>27</b>
6.1	Measurement methodology . . . . .	27
6.1.1	Booter attacks . . . . .	28

<b>7 Making the Case for Elliptic Curves in DNSSEC</b>	<b>30</b>
7.1 Deployment Scenarios . . . . .	30
7.2 Fragmentation . . . . .	31
7.2.1 The problem . . . . .	31
7.2.2 Revisiting Fragmentation using Elliptic Curves . . . . .	31
<b>8 Mitigation of Topological Inconsistency Attacks in RPL based Low Power Lossy Networks</b>	<b>33</b>
8.1 Introduction . . . . .	33
8.1.1 Context . . . . .	33
8.1.2 The RPL protocol . . . . .	33
8.1.3 Attack description . . . . .	33
8.2 Results on attack mitigation and detection . . . . .	34
8.2.1 Mitigation approaches . . . . .	34
8.2.2 Results . . . . .	35
<b>9 Impact of Packet Sampling on Link Dimensioning</b>	<b>36</b>
9.1 Research Problem . . . . .	36
9.2 Contributions . . . . .	36
9.3 Proposed Approaches . . . . .	36
9.4 Results and Validation . . . . .	38
9.5 Concluding Remarks . . . . .	39
<b>10 Integration of EU research</b>	<b>40</b>
10.1 International Activities . . . . .	40
10.2 Collaborations with Other EU Projects and Institutions . . . . .	41
<b>11 Conclusions</b>	<b>42</b>
<b>A Internet of Names - Poster</b>	<b>45</b>
<b>B Determining the State of Security in the IPv6 Internet</b>	<b>47</b>
<b>C Characterizing and Mitigating the DDoS as a Service Phenomenon</b>	<b>49</b>
<b>D DDoS Attack Mitigation using OpenFlow-based SDN</b>	<b>51</b>

# 1 Introduction

Network and service monitoring is at the basis of any informed management decision. As such, monitoring is one of the cornerstones of the management of the Future Internet, and one of the core research areas of FLAMINGO. In Y3, WP5 has continued the positive trend in research in the field of network and service monitoring that was shown in Y1 and Y2, with an excellent scientific output. The goal of this deliverable is to describe FLAMINGO's achievements in this research domain.

The deliverable is structured such as to give relevance to the objectives set for this WP. Section 2 reports the S.M.A.R.T. (Specific, Measurable, Achievable, Relevant, Timely) objectives and how the WP has successfully achieved them. It then summarizes the active research that it is taking place on the topics identified by the WP5-specific objectives.

The first S.M.A.R.T. objective is the *integration of PhD students*. In adherence to the Description of Work (DoW), at least two fully integrated PhD students are active in WP5. In addition, several PhD collaborations have been active during Y3. Information regarding these topics can be found in Section 3.

The second S.M.A.R.T. objective concerns the *scientific output* of the project. In Y3, the research work packages have published 73 papers, at major conferences and in journals. Many of the published papers have a measurement component that has been investigated in the context of WP5. A mapping of the papers to the relevant WPs is presented in D8.3. Also in Y3, WP5 has targeted high-ranking publication venues in the areas of networking, network monitoring as well as network management, which resulted in publications in conferences IMC 2014 and SIGCOMM 2015 (short paper), and in journals such as ACM Computer Communication Review and Transactions on Network and Service Management (see also Sections 4–9). An overview of the status of the S.M.A.R.T. objective is given in Section 2.1. For a detailed list of the FLAMINGO published papers we refer the reader to D8.3.

Sections 4–9 highlight a selection of the current research activities that are related to the WP5-specific objectives. In line with a trend that already emerged in Y1 and Y2, also Y3 has indicated that large scale measurements and measurements of security events are two of the core areas of interest for WP5. The highlights of research focus therefore on these topics. In particular, Section 4 reports on the lesson learned in extensively using the distributed measurement platform RIPE Atlas for network measurements; Section 5 describes a new measurement infrastructure that has been developed at UT in collaboration with SURFnet, aiming at collecting daily snapshot of the state of DNS by querying all domain names registered in .com, .net and .org zones (roughly equivalent to 50% of the registered domain names worldwide); Section 6 proposes a characterization of Booters, online webservices providing DDoS-as-a-Service; Section 7 advocates, based on evidences supported by network measurements, the use of elliptic curve cryptography in place of RSA in DNSSEC; Section 8 investigates the detection and mitigation of topological inconsistency in the Routing Protocol for Low-power Lossy Networks (RPL); finally, Section 9 structurally investigates the impact of packet sampling on link dimensioning.

## 2 Objectives and Tasks

This section presents an overview of the S.M.A.R.T. (Specific, Measurable, Achievable, Relevant, Timely) objectives for WP5 and the WP5-specific objectives. For the S.M.A.R.T. objectives, we indicate how these have been achieved for Y3. For the WP5-specific objectives, we summarise the activities that have taken place among the consortium partners, and we indicate the plans to address some of those objectives in subsequent work.

### 2.1 S.M.A.R.T. Objectives

To comply to the S.M.A.R.T. objectives, WP5 has been active on the following topics:

- **Integration of Ph.D. students** – The Description of Work (Section B.1.1.5) states that “after 9 months each research WP will have identified at least two fully integrated Ph.D. students, which means that these students will be jointly supervised and financially paid by FLAMINGO”. In the first two years of the project 14 PhD students have joined FLAMINGO as *fully integrated PhD students*. In the third year, three more PhD students have joined the NoE. These students, their affiliations and the co-supervising institutions are listed in D8.3. Collaborations are a cornerstone of research within FLAMINGO. It is important to note that collaborations are not only taking place between fully integrated PhD students, but also among students that are not financially paid by FLAMINGO but that are actively contributing to the WP work. More information on the integration of Ph.D. students, the Ph.D. students active in the context of WP5 and their collaborations within the consortium can be found in Section 3, which report the description of the collaborations that are more closely related to the work in this WP.
- **Research** – The Description of Work (Section B.1.1.5) states that “after 18 month at least 20 scientific papers will be submitted / published”. In the first two years the project had exceeded the expected number of publications. In Y3 the research work packages published 73 papers at major conferences as well as in journals, and with this exceed the expected number of papers. In addition, several other papers are currently under review. The complete list of published papers and a mapping of the papers to the relevant WPs can be found in D8.3. We also highlight in this deliverable the scientific output of WP5 with respect to the collaboration with other EU projects and institutions and within the FLAMINGO consortium (Tables 1–3). Last, the partners have also targeted, together with the top conferences and journals in the network management field, high-end conferences and journals in the wider field of networking and network measurements. This effort resulted in a paper and an editorial in the ACM SIGCOMM Computer Communication Review (CCR), a paper published at Internet Measurement Conference (IMC 2014), a short paper (poster) at SIGCOMM 2015.

### 2.2 Workpackage Objectives

This section provides a high-level summary of the WP5-specific objectives (as identified in the DoW, WP5 description table). These objectives have been grouped into two categories. Section 2.2.1 describes the status of the objectives in which WP5 researchers are currently active, both in terms of research topics as well as academic activities in general. We refer to these as *ongoing objectives*. Section 2.2.2 includes the objectives for which no activity is currently being carried out in this WP (*open objectives*).

Table 1: WP5 publications in collaboration with other EU projects and institutions.

Authors	Title	Venue	EU project/ institution
V. Bajpai, S.J. Eravuchira, and J. Schönwälder	Lessons Learned from using the RIPE Atlas Platform for Measurement Research	ACM/SIGCOMM CCR	Leone
M. Jonker and A. Sperotto	Mitigating DDoS Attacks Using OpenFlow-Based Software Defined Networking	AIMS 2015	NWO Project D3
L. Hendriks, A. Sperotto, and A. Pras	Characterizing the IPv6 Security Landscape by Large-Scale Measurements	AIMS 2015	SURFnet
J.J. Santanna, R. Durban, A. Sperotto, and A. Pras.	Inside booters: An analysis on operational databases	IM 2015	INSA of Toulouse
J.J. Santanna, R. van Rijswijk-Deij, R. Hofstede, A. Sperotto, M. Wier- bosch, L. Zambenedetti Granville, and A. Pras	Booters – An analysis of DDoS-as-a-service attacks	IM 2015	SURFnet, UFRGS
O. van der Toorn, R. Hofstede, M. Jonker, and A. Sperotto	A first look at HTTP(S) intrusion detection using NetFlow/IPFIX	IM 2015	SALUS
D. van der Steeg, R. Hofstede, A. Sperotto, and A. Pras	Real-time DDoS attack detection for Cisco IOS using NetFlow	IM 2015	SALUS
M. Jonker, R. Hofstede, A. Sperotto, and A. Pras	Unveiling flat traffic on the Internet: An SSH attack case study	IM 2015	SALUS
R. van Rijswijk-Deij, A. Sperotto, and A. Pras	DNSSEC and Its Potential for DDoS Attacks: A Comprehensive Measurement Study	IMC 2014	SURFnet
G. Machado, T. Bocek, A. Filitz, and B. Stiller	Measuring interactivity and geographical closeness of online social network users in support of social recommendation systems	CNSM 2014	SmartenIT
A. Lareida, T. Bocek, M. Pernebayer, and B. Stiller	Automatic network configuration with dynamic churn prediction	IM 2015	SmartenIT

### 2.2.1 Ongoing Objectives

**Objective 1: To integrate European research in the area of (flow-based) network and service monitoring** – In collaboration with WP3, WP5 has taken part in several activities related to network and service monitoring at the European level. An example is the upcoming Dagstuhl seminar *Global Measurements: Practice and Experience* (January 4-7, 2016), in which FLAMINGO members are active in the organization (J. Schönwälder) and as par-

Table 2: FLAMINGO publications in collaboration with other EU projects and institutions (ctd.).

Authors	Title	Venue	EU project/ institution
C. Tsiaras, M. Rösch, and B. Stiller	VoIP-based Calibration of the DQX Model	Networking 2015	SmartenIT
D.Dönni, G.S.Machado, C.Tsiaras, and B.Stiller	Schengen Routing: A Compliance Analysis	AIMS 2015	SmartenIT
T. Bocek, N. Rutishause, and B. Stiller	Energy-efficient Overlay Networks for Mobile Devices with Buffered Relaying and Push Notifications	LCN 2015	SmartenIT
N. Bouten, R. de O. Schmidt, J. Famaey, S. Latré, A. Pras, and F. De Turck	QoE-driven in-network optimization for Adaptive Video Streaming based on packet sampling measurements	Computer Networks	University of Antwerp, iMinds V-FORCE
J. van der Hooft, S. Petrangeli, M. Claeys, J. Famaey, and F. De Turck	A Learning-Based Algorithm for Improved Bandwidth-Awareness of Adaptive Streaming Clients	IM 2015	University of Antwerp
R. van Rijswijk-Deij, M. Jonker, A. Sperotto, and A. Pras	The Internet of Names: A DNS Big Dataset	SIGCOMM 2015 (short paper)	SURFnet
R. van Rijswijk-Deij, A. Sperotto, and A. Pras	Making the Case for Elliptic Curves in DNSSEC	ACM/SIGCOMM CCR	SURFnet
R. De O. Schmidt, R. Sadre, A. Sperotto, H. van den Berg, and A. Pras	Impact of Packet Sampling on Link Dimensioning	TNSM	Universef
R. De O. Schmidt, H. van den Berg, and A. Pras	Measurement-based network link dimensioning	IM 2015	Universef, MCN, SURFnet Gigaport
R. De O. Schmidt, L. Hendriks, A. Pras, and R. van der Pol	OpenFlow-Based Link Dimensioning	SC 2014	MCN, SURFnet Gigaport

ticipants. WP5 has also organized the IJNM special issue “Measure, Detect and Mitigate Challenges and Trends in Network Security”, a collaboration between UT, UZH, UniBwM and CAIDA. Last, another important event has been the organization of the 9th International Conference on Autonomous Infrastructure, Management and Security (AIMS 2015). Besides having several FLAMINGO partners involved in the organization of the conference, we highlight also the WP-5 related Lab sessions “Map-Reduce and Hadoop” by J. François (INRIA) and “Powering Monitoring Analytics with ELK Stack” by A. Lahmadi and F. Beck (INRIA). For additional information about activities related to this objective, we refer the reader to Section 10 and Deliverable D3.4.

**Objective 2: To create and maintain articles within Wikipedia and other online systems in this area** – In Y3, WP5 has collaborated in the maintaining of a set of Wikipedia pages

Table 3: WP5 Publications authored by multiple FLAMINGO partners.

Authors	Title	Venue	FLAMINGO partners
N. Bouten, R. de O. Schmidt, J. Famaey, S. Latré, A. Pras, and F. De Turck	QoE-driven in-network optimization for Adaptive Video Streaming based on packet sampling measurements	Computer Networks	iMinds, UT
A. Mayzaud, A. Sehgal, R. Badonnel, I. Chrismont, and J. Schj�onw�lder	Mitigation of Topological Inconsistency Attacks in RPL based Low Power Lossy Networks	IJNM	INRIA, JUB

related to networking: *Software Defined Networking*<sup>1</sup>, *Schengen Routing*<sup>2</sup>, *Content delivery network*<sup>3</sup>.

**Objective 5: To collect (anonymized) monitoring data** Also in Y3, WP5 has been involved in several measurements activities. Some of those have continued from Y2. Others have instead started in Y3. The *Internet Traffic Statistics* project, introduced in Y1 and further developed in Y2, is currently collecting weekly traffic reports from supporting operators (the Brazilian NREN RNP, the Czech NREN Cesnet, the Danish NREN DeiC). The projects is currently under active development such as to provide easy access to data for users and easy data retrieval from supporting partners.

Several measurement activities fall into the category of active measurements. Notable examples are the Booter measurements conducted at the University of Twente in collaboration with SURFnet, with the aim of characterising and fingerprinting DDoS attacks. Additional information can be found in Sec. 6. The work in [2] and [3] both make use of RIPE Atlas probes for topological measurements. Both papers have been awarded a **best paper award** at AIMS 2015. Finally, in Y3, the University of Twente, in collaboration with SURFnet and SIDN Labs<sup>4</sup> (the organization managing the .nl DNS zone), has started a large scale DNS measurement aimed at systematically querying domain names worldwide, with the goal of progressively building a reliable picture of the status of DNS over time. Currently, the project is querying, daily, the .com, .org and .net domain names (equivalent to around 50% of the registered domain names worldwide). A short paper detailing the measurement has been accepted at SIGCOMM 2015 [4]. For more information, see Sec. 5.

Finally, monitoring data for specific application domains, in this case Internet of Things, have also been carried on in the context of the research on security aspects on RPL (see Sec. 8).

More information about the data collection activities carried on in WP5 can be found in D1.3.

**Objective 6: To create annotated traces to assess the quality of different Intrusion Detection Systems** – Several activities that aimed at creating annotated traces for Intrusion Detection have been carried on in Y3. A notable contribution to this objective is given by the research conducted at the UT in the context of SSH and HTTP(S) flow based intrusion detection, which was started in Y2, continued in Y3 and concretized in several publications.

<sup>1</sup>[https://en.wikipedia.org/wiki/Software-defined\\_networking](https://en.wikipedia.org/wiki/Software-defined_networking)

<sup>2</sup>[https://en.wikipedia.org/wiki/Schengen\\_Routing](https://en.wikipedia.org/wiki/Schengen_Routing)

<sup>3</sup>[https://en.wikipedia.org/wiki/Content\\_delivery\\_network](https://en.wikipedia.org/wiki/Content_delivery_network)

<sup>4</sup><https://www.sidnlabs.nl/>

UT has extended the research started in Y2 on developing a targeted detection scheme for HTTP(S)-based brute force attacks (targeting web applications with the intention of gaining unauthorized access to password-protected pages). The original work, conducted in the context of the B.Sc. thesis of O. van der Toorn (supervised by R. Hofstede) aimed at analyzing malicious code attack patterns against commonly used content management systems (e.g., Joomla, Wordpress and Drupal) and at extracting appropriate flow-level signatures. Such a work has been published in Y3 in IM 2015 [5]. The limitation of this work, however, was that the devised detection mechanism was not able to fully capture the dynamic behavior of HTTP(S) attacks. To overcome this problem, a new approach has been investigated, this time based on additional flow-based information. In particular, leveraging on the flexibility of IPFIX, a dedicated flow definition exporting histogram of payload sized for the observed packets has been used to gather additional information about the attack traffic. A clustering mechanism has then be applied to separate benign from malicious traffic. The paper is currently under submission at INFOCOM 2016.

In addition, we report that the work initiated in Y2 on the analysis of the effect of TCP retransmissions and control information on the detection of SSH dictionary attacks (validated using the open source tool SSHCure) has been accepted at IM 2015 [6].

**Objective 7: To investigate the applicability of different AI and machine learning techniques for flow analysis** – In Y3, several activities on the topic of machine learning and its application to network traffic has been carried out.

One of those is the work in [7] which applies Reinforcement Learning to the problem of optimizing client parameter information in the context of HTTP Adaptive Streaming (HAS). The work has been published in IM 2015.

Machine learning is also the approach investigated in the work of the PhD student C. Dietz (UniBwM and UT), who aims at applying ML techniques to the detection of Botnets. The currently ongoing work aims at identifying the most informative features, able therefore to capture Botnet behavior. The analysis of the feature is based on sink-hole data collected for the Kelihos botnet, one variant of which as been sink-holed in 2013 by a team of security expert among which SURFnet.

Finally, clustering has been applied to the HTTP(S) dictionary attack detection investigated at UT and currently under submission at INFOCOM 2016 (see also Objective 6).

**Objective 8: To propose novel solutions for intrusion detection and fingerprinting** – Research in the field of intrusion detection and fingerprinting in Y3 has focused on the topics of Booter blacklisting and collaborative mitigation.

Distributed Denial of Service (DDoS) attacks mean millions in revenue losses to many industries, such e-commerce and online financial services. The amount of reported DDoS attacks has increased with 47% compared to 2013. One of the reasons for this increase is the availability and ease of accessibility to websites, which provide DDoS attacks as a paid service, called Booters. Although there are hundreds of Booters available, current researches are focused on a handful sample of them - either to analyse attack traffic or hacked databases. Towards a thorough understanding and mitigation of Booters, a comprehensive list of them is needed. The work in [8] characterizes Booter websites and demonstrates that the identified main characteristics can be used to classify Booters with 85% of accuracy. The outcome of the classification is a blacklist of Booters, that can be used to either filter or notify tentative access to Booters website. The generated blacklist is currently being tested by the Dutch National Research and Education Network (SURFnet).



A different initiative is the one carried on by UT in collaboration with the University of Applied Science Darmstadt (PhD student J. Steinberger), where possible schemes for collaborative intrusion detection are investigated. The rationale behind this research stems from the observation that, for certain categories of attacks, e.g., DDoS attacks, mitigation close to the target is often non-effective, mostly due, for example to resource exhaustion. To overcome this issue, mitigation should take place closer to the network backbone, when the volume of traffic is still manageable. However, this implies that several parties need to take co-ordinated action in order to successfully contain an attack. How coordination could work and how effective it could be in practice is the focus of this research. This research is conducted in collaboration between WP5 and WP6.

**Objective 9: To propose and study monitoring frameworks for IaaS, PaaS and SaaS Clouds (i.e., to allow elastic management of cloud infrastructures)** – INRIA (G. Hurel) is investigating the integration of security functions in clouds in the context of mobile devices. Security applications may have a significant impact on the device resources. Users may be tempted to uninstall or disable them with the objective of increasing battery lifetime and avoiding configuration operations and updates. To overcome these issues, INRIA proposed an approach based on outsourcing mobile security functions and building transparent in-path security compositions for mobile devices. The outsourced functions are dynamically activated, configured and composed using software-defined networking and virtualization capabilities. The work has been published in IM 2015 [9].

## 2.2.2 Open Objectives

**Objective 4: To develop a flow query language for expressing temporal relationships of complex flow patterns** This objective has been achieved in Y1, with the work carried out at JUB on the *Network Flow Query Language* (NFQL). However, since the tool related to the NFQL research is still being developed as part of WP1, more activities might be linked to this objective in the future. For this reason, we now report it as “open”.

**Objective 3: To develop a generic distributed flow monitoring architecture** – In Y1, WP5 and WP6 have proposed a monitoring architecture that provides a consistent view of the FLAMINGO efforts in network and service monitoring and on configuration and repair. For WP5, the development of this architecture has been completed in Y1, and in Y2 and Y3 the architecture development has been carried out by WP6, while WP5 has supported the architecture when necessary by providing support in terms of measurements. For this reason, the objective is reported as “open”. For more details about the architecture, we refer the reader to D6.3.

## 2.3 Tasks and Objectives Mapping

Table 4 summarises the status of the S.M.A.R.T. objectives related to WP5 (Section 2.1) and the WP5-specific objectives (Section 2.2). For each of the considered objectives, Table 4 indicates if the objective has been achieved (S.M.A.R.T objectives), or if there are WP activities that are contributing to the objective (WP5-specific objectives). For the WP5-specific objectives, Table 4 shows to which of the Tasks in the DoW the objective is contributing to. Finally, the table acts as a guide for the reader to locate the sections of this deliverable that provide additional information on a specific objective.

Table 4: Mapping of objectives and tasks.

<b>Objective</b>	<b>Task 5.1</b>	<b>Task 5.2</b>	<b>Task 5.3</b>	<b>Status</b>	<b>Additional Material</b>
S.M.A.R.T. Objective 1 Integration Ph.Ds				Achieved	Section 3
S.M.A.R.T. Objective 2 Research				Achieved	D8.2
WP Objective 1 Integration EU Research				Ongoing	Section 10
WP Objective 2 Articles Online Systems				Ongoing	D2.3
WP Objective 3 Monitoring Architecture	X			Open	D6.3
WP Objective 4 Flow Query Language			X	Open	
WP Objective 5 Monitoring Data		X		Ongoing	D1.3 (data collection) Section 4 Section 5 Section 7 Section 9
WP Objective 6 IDS and traces		X		Ongoing	
WP Objective 7 AI & Machine Learning		X	X	Ongoing	
WP Objective 8 ID & Fingerprinting			X	Ongoing	Section 6
WP Objective 9 Cloud Infrastructures			X	Ongoing	

Table 5 indicates the progress WP5 has made with respect to the S.M.A.R.T. objectives and the WP5-specific objectives. For the WP5-specific objectives, which are not directly measurable, we report the keywords of the core activities conducted for the objective.

Table 5: Progress with respect to Y2.

<b>Objective</b>	<b>Y2 activities</b>	<b>Y3 activities</b>
S.M.A.R.T. Objective 1 Integration Ph.Ds	14 Integrated Ph.D. (total)	17 Integrated Ph.D. (total)
S.M.A.R.T. Objective 2 Research	50 papers (entire project)	73 papers (entire project)
WP Objective 1 Integration EU Research	NMRG; Dagstuhl; AIMS; Coll. EU level	NMRG; Dagstuhl; AIMS; Coll. EU level
WP Objective 2 Articles Online Systems	Wikipedia contribution	Wikipedia contribution
WP Objective 3 Monitoring Architecture	Support for WP6	Support for WP6
WP Objective 4 Flow Query Language	Not Addressed	Not Addressed
WP Objective 5 Monitoring Data	Longitudinal analysis DNS; DNSSEC; SSH Link Dimensioning	DNS; DDoS attacks; packet sampling
WP Objective 6 IDS and traces	Compromise detection SSH Detection HTTP(S) detection	SSH and HTTP(S) detection; IPv6 security
WP Objective 7 AI & Machine Learning	HTTP adaptive streaming Botnet detection HTTP(S) dictionary attack detection	HTTP adaptive streaming Botnet detection
WP Objective 8 ID & Fingerprinting	SSH Detection HTTP(S) detection	SSH Detection HTTP(S) detection Booter traces
WP Objective 9 Cloud Infrastructures	Security and Clouds Cloud Networking	Security in Clouds

### 3 Integration of PhD Students

The integration of PhD students is one of the S.M.A.R.T. objectives within this WP. Section 3.1 gives an overview of FLAMINGO collaborations that are ongoing, ended or are in the process of starting during this year. Furthermore, collaborations envisioned for the next year of FLAMINGO are shown.

In the FLAMINGO approach, monitoring (WP5) is at the basis of any configuration and repair action (WP6), while both activities are conducted within the boundaries of the economic, legal and regulative constraints (WP7). Y3 has seen several collaborations between the three research WPs, and in particular between WP5 and WP6. However, in answer to the reviewers comments, please note that the PhD collaborations have now been mapped to the WP that most closely relate the them. This deliverable, therefore, reports the subset of PhD collaborations that are closer to the research conducted in WP5 (Section 3.2). For WP6 PhD collaborations, we refer the reader to D6.3. Finally, please note that all fully integrated PhD students are listed in D8.3, including their co-supervisors and affiliation.

Table 6: Overview of the FLAMINGO Collaborations, as in Figure 1

Acronym	Researchers	WPs	Status
INRIA-JUB-RPL	A. Mayzaud - A. Sehgal	WP5, WP6	Ended
INRIA-JUB-Distr	A. Mayzaud - A. Sehgal	WP5, WP6	Ended
INRIA-UniBwM-Cloud	S. Seeber - G. Hurel	WP6	Ongoing
UCL-iMinds-Cache	D. Tuncer - M. Claeys	WP5, WP6	Ongoing
UT-UniBwM-IDS	R. Hofstede - M. Golling	WP5, WP6	Ongoing
UT-INRIA-Flowoid	R. Hofstede - A. Lahmadi	WP5	Ongoing
UZH-UniBwM-JUB-M2	C. Tsiaras - S. Seeber D. Doenni - A. Sehgal	WP5, WP7	Ended
iMinds-UPC-NetVirt	N. Bouten - R. Mijumbi M. Claeys	WP6	Ongoing
INRIA-UniBwM-Chain	G. Hurel - S. Seeber	WP6	Ongoing
UZH-UniBwM-UT-Scheng	C. Tsiaras - M. Jonker S. Seeber - L. Stiemert	WP5, WP6	Ongoing
UniBwM-UT-MLB	C. Dietz - A. Sperotto	WP5, WP6	Ongoing
UT-UniBwM-Class	J. Santanna - C. Dietz	WP5	Planned
UT-UniBwM-Sec	A. Pras - S. Seeber	WP5, WP6	Starting
iMinds-UT-OpenFlow	S. Petrangeli, R. Schmidt	WP5, WP6	Planned
JUB-UT-Pattern	N. Melnikov - R. Schmidt	WP5, WP6	Ended
UT-JUB-Booters	J. Santanna - A. Sehgal	WP5	Ended
UniBwM-JUB-RPL	S. Seeber - B. Stelte - A. Sehgal	WP6	Ended
UCL-UT-MAN	D. Tuncer - R. Schmidt	WP5, WP6	Ended
iMinds-UT-QoS	R. Schmidt - N. Bouten	WP5, WP6	Ended

### 3.1 PhD Student Collaborations

The integration of PhDs into FLAMINGO enabled valuable and fruitful joint research in the area of network and service management. The bottom-up approach in the previous years was continued to integrate experienced researchers as well as new researchers not necessarily paid by FLAMINGO. Table 6 summarizes the collaborations, the affiliations involved and their respective status. Each collaboration can have one of the following status: **ONGOING**, **ENDED**, **STARTED**, **PLANNED**. **ENDED** applies to collaborations started in Y1 and Y2 of FLAMINGO and ended in Y3 because the research goals have been reached. A collaboration is called **ONGOING** if started during Y3 or previous years and progress is already reported (e.g. measurement results, planned papers, ...). **STARTING** collaborations are in the process of defining their topic, research interests and goal of the collaboration, and drafting a plan how to possibly reach their goal. The last type of collaborations with the status **PLANNED** have defined mutual interest in working jointly together, but didn't define a concrete topic.

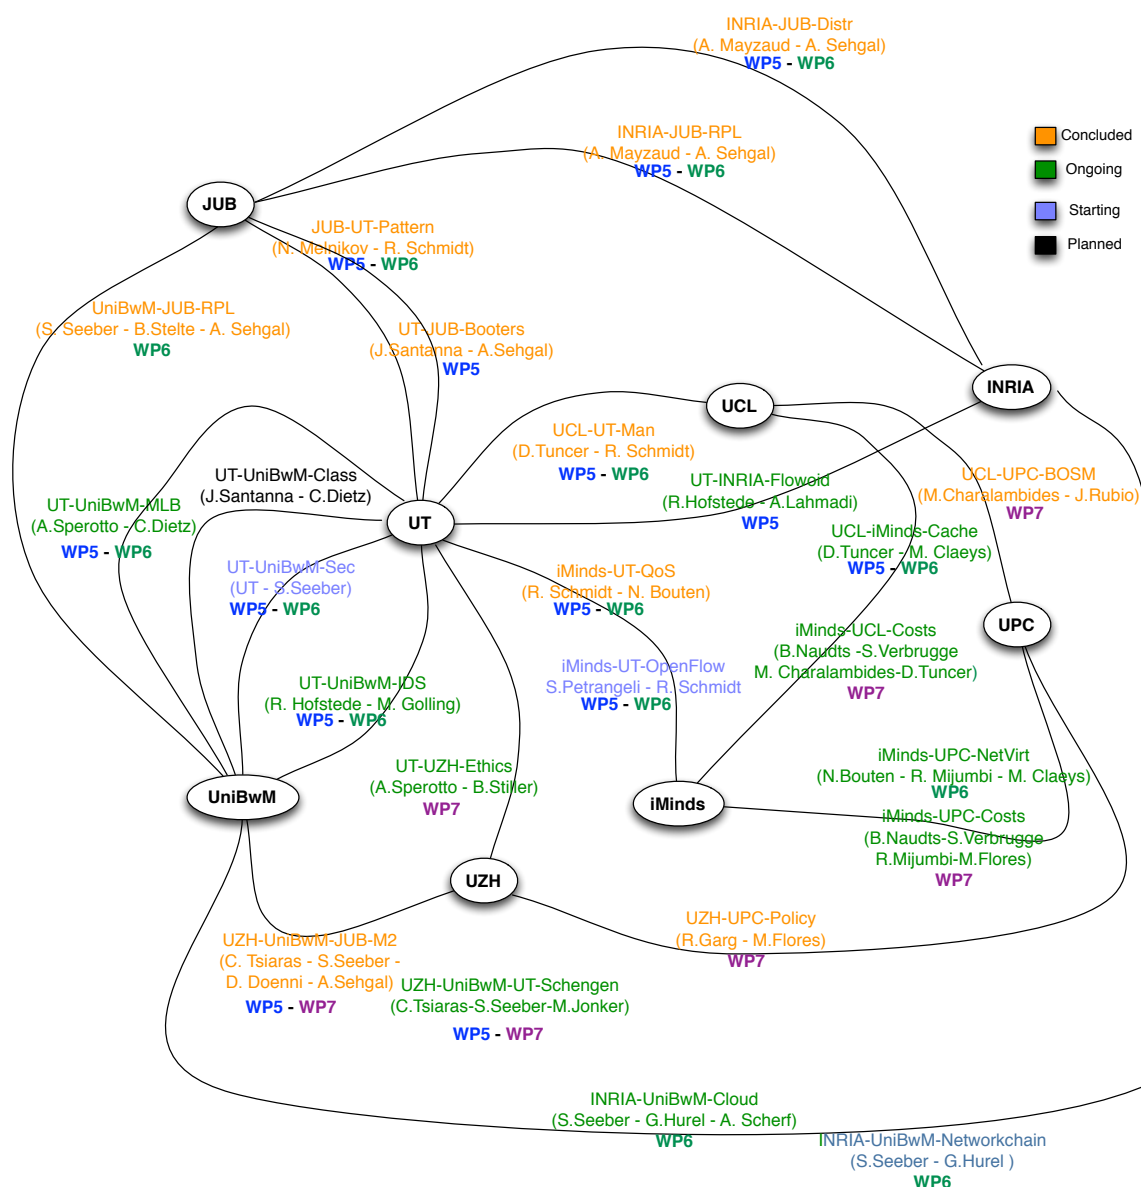


Figure 1: Overview of PhD collaborations

### 3.2 Description of the Collaborations

This section presents the currently ongoing and recently ended collaborations that are more closely related to WP5. For WP6-related collaborations, we refer the reader to D6.3. Each collaboration description roughly follows the same structure. At first the topic of each collaboration is explained. Subsequently the progress and achievements in Y3 are highlighted. Depending on the status of a collaboration further steps are described. In the end each collaboration highlights the contribution WP5. Table 7 reports the PhD students that are involved in the PhD collaborations reported by WP5 and WP6.

### 3.2.1 Intrusion Detection Systems (UT-UniBwM-IDS)

This joint research activity is a collaboration between UT and UniBwM. Intrusion detection is nowadays commonly performed in an automated fashion by IDSes [10]. Several classifications for IDSs

Table 7: PhD students involved in Ongoing WP5/WP6 collaborations

Name	Affiliation	Collaborations	Acronym
Anth��a Mayzaud	INRIA	JUB	INRIA-JUB-RPL INRIA-JUB-Distr
Gaetan Hurel	INRIA	UniBwM	INRIA-UniBwM-Cloud
Rick Hofstede	UT	UniBwM, INRIA	UT-INRIA-Flowid UT-UniBwM-IDS
Ricardo Schmidt	UT	UCL	UCL-UT-Man
Mario Golling	UniBwM	UT	UT-UniBwM-IDS
Sebastian Seeber	UniBwM	UZH, JUB, INRIA	UZH-UniBwM-JUB-M2 INRIA-UniBwM-Cloud INRIA-UniBwM-Chain UT-UniBwM-Scheng
Rashid Mijumbi	UPC	iMinds	iMinds-UPC-NetVirt
Anuj Sehgal	JUB	INRIA, UT, UniBwM	INRIA-JUB-RPL INRIA-JUB-Distr UZH-UniBwM-JUB-M2
Christos Tsiaras	UZH	UniBwM, JUB	UZH-UniBwM-JUB-M2
Daniel D��nni	UZH	UniBwM, JUB	UZH-UniBwM-JUB-M2
Niels Bouten	iMinds	UPC	iMinds-UPC-NetVirt
Maxim Claeys	iMinds	UCL	iMinds-UPC-NetVirt UCL-iMinds-Cache
Mattijs Jonker	UT	UniBwM	UT-UniBwM-Scheng
Christian Dietz	UniBwM	UT	UniBwM-UT-MLB

are common. One of these classifications focuses on the kind of data that is used for performing intrusion detection. The first class of IDSs mainly uses packet headers (flows) for intrusion detection. While these flow-based IDSs have a high-performance and are usually little privacy-intrusive, they are typically affected by a high number of undetected attacks (false negatives; see Figure 2). In contrast to flow-based IDSs, payload-based IDSs are capable of performing extensive layer-7-detection (and, therefore, have a lower false negative rate), but at the expense of a much higher system requirements as well as a violation of privacy [11].

Given these observations, performing intrusion detection in high-speed networks is a challenging task. While many payload-based IDSs are working well at the backend of service provider networks, the backbone is often characterized by communication links with high-speed connections and thus requires well equipped IDS in order to be capable of handling 100 Gbps or more, for example [12]. Within this collaboration, it is planned to create a framework for distributed intrusion detection in high-speed networks by combining especially flow-based and payload-based intrusion detection. As already stated, in addition to monetary aspects, legal issues in general and privacy issues in particular are also important reasons why payload-based IDS are rarely deployed in high-speed networks today [11].

In order to overcome these disadvantages, this collaboration makes use of both approaches (flow-based and payload-based intrusion detection) in a multi-layered approach. As depicted in Figure 3, the approach is centered around the ideas that (i) the first detection layer comprises flow-based intrusion detection, which performs detection based on the entire packet stream (100%) and that (ii) depending on the result of the flow-based detection, the payload-based IDS is activated for a certain period of time to investigate the anomaly of the flow-based IDS in more detail (1%) – in order to verify or falsify the result of the flow-based IDS. As network attacks can last shortly and a

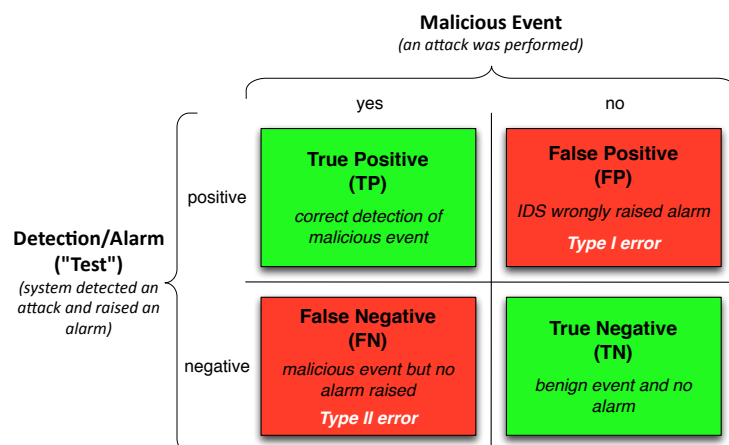


Figure 2: Categories of Alarms ("Confusion Matrix").

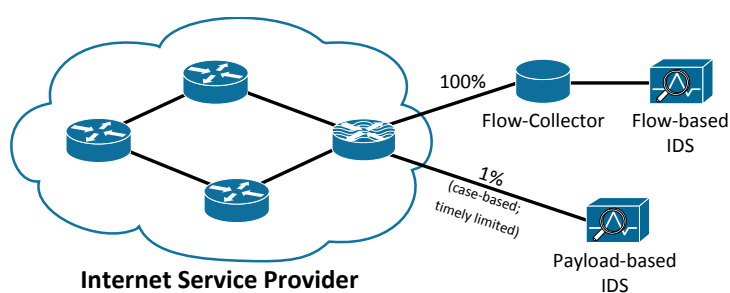


Figure 3: Simplified scenario for UT-UniBwM-IDS.

switch has to be made from flow-based to packet-based detection, detection has to be performed in real-time.

All the ideas presented above are summarized in our architecture for multi-layered intrusion detection, shown in Figure 4, and published in [13]. The architecture features three main data streams:

- A Flow meta-data that can be retrieved directly from the router's Command-Line Interface (CLI).
- B Flow data, exported by means of Cisco's NetFlow [20] or the recent IETF standardization effort IPFIX.
- C Full packet streams, potentially pre-filtered by the router upon instruction by the *Manager*.

Key characteristic of the *Real-Time IDS* is that it constantly analyses the full traffic stream, without any form of sampling or filtering. In a previous work, we have shown that a similar approach is able to mitigate DDoS attacks in near real-time [14]. Upon detection of such an attack, the *Real-Time IDS* can reconfigure the router to drop the attack traffic, to make sure that neither the network itself, nor the monitoring infrastructure is overloaded. In addition, the *Manager* is informed about the attack by means of a standardized message exchange format, such as the Intrusion Detection Message Exchange Format (IDMEF).

Besides the *Real-Time IDS*, also the *Flow-Based IDS* is constantly monitoring its input data stream. Given that flow export technologies, such as NetFlow and IPFIX, aggregate packets into flows, such an IDS is usually capable of monitoring the aggregated traffic using commodity hardware. An example of a flow-based IDS is SSHCure,<sup>5</sup> which detects SSH dictionary attacks and reports whether a host has been compromised [15]. The *Flow-Based IDS* may be informed by the *Manager* about previous detections, and reports its own detections to the *Manager* again. Although not supported by current IDSs, the main idea of forwarding previous detection results to IDSs is to give as much information as possible and so to make the process of intrusion detection as reliable as possible.

In situations where the *Manager* decides to initiate a more extensive analysis of an attack, the *Protocol-Based IDS* or *DPI-based IDS* can be activated and instructed. The *Manager* decides which IDS is most suitable for a particular attack. Before activating the other IDSs, the *Manager* has to reconfigure the router to pre-filter the traffic stream to only include the attack traffic. Analogously to the *Flow-Based IDS*, these IDSs report their detections to the *Manager*. If an attack has been detected, the router is instructed to drop the attack traffic. If an attack could not be confirmed, the *Manager* will not dispatch any investigation about that particular traffic to the various IDSs anymore.

As to the state of the multi-layered architecture, we can report that a prototype has been developed that is currently under validation. We are currently in the process of collecting the first datasets, which will be used for testing the operation of the prototype and, in a later stage, for performing a validation of the accuracy.

The collaborative work contributes mainly to WP5: Network and Service Monitoring by performing multi-layered IDS. Especially objective 8 - novel solutions for IDS is addressed with this collaboration. As the long-term goal of this collaboration is also to link different managers with each other, this addresses objective 3 "to develop a generic distributed flow monitoring architecture". Regarding this objective, in [16] an Evaluation of State of the Art IDS-Message Exchange Protocols was already performed.

---

<sup>5</sup><http://github.com/sshcure/sshcure/>






Figure 4: Architecture for multi-layered intrusion detection [13].

### 3.2.2 Flowoid: a NetFlow/IPFIX Probe for Android-based Devices (UT-INRIA-Flowoid)

Analysis of the network behaviour of applications running on a smartphone device requires the collection of information about data leaving the device and where it is sent. Cisco's NetFlow and the more recent IPFIX are flow export technologies that have seen a rapid adoption and widespread integration in many campus, enterprise and backbone networks. To be able to export and analyse mobile device characteristics (such as its location at the moment of certain network activity), the NetFlow and IPFIX protocols have to be extended. The flow exporter, flow collector and analysis application need to be aware of these extensions as well.

The work in this collaboration has been divided between INRIA and UT. On the one hand, INRIA is responsible for developing a flow exporter tailored to Android devices. On the other hand, UT is responsible for the flow collector and analysis application. The major achievements of the collaboration are:

- Development of a NetFlow and IPFIX metering process for Android devices;
- Extension of nfdump/Nfsen and SURFmap with location support;
- IETF Internet-Draft (ID) describing a set of information elements for IPFIX metering process location [17].

The following aspects of this collaboration fall within WP5. In this work, INRIA has developed Flowoid, a NetFlow and IPFIX metering process tailored to Android devices. The probe associates geolocation data with each observed network flow, consisting of the GPS coordinates of the mobile device, among others. This information is exported together with the traditional fields defined in the NetFlow and IPFIX: IP version, source and destination addresses, the number of exchanged bytes, the type of protocol, the number of exchanged packets, the source and destination ports, and the

duration of a flow. In addition, it contains seven additional fields that denote the identifier of the device: the identifier of the localization method, a timestamp, the integer part of the latitude, the decimal part of the latitude, the integer part of the longitude, and the decimal part of the longitude. UT has extended the state-of-the-art flow collector nfdump/NfSen<sup>6</sup> with location support, allowing us to analyse the flows exported by the Flowoid probe. In a previous work UT has developed a network monitoring tool based on the Google Maps API, named SURFmap [18],<sup>7</sup> which adds a geographical dimension to flow data and displays the data on a map. Since SURFmap [18] already supports network traffic geolocation (i.e. adding physical locations of hosts to network data), this tool has been extended to visualize the locations of devices on a map. This will allow us to visualize network traffic of mobile device with respect to devices locations.

The main activity carried out in Y3 is related to improving the Internet Draft (ID) ([17]) based on community feedback. Details regarding the progress in Y3 on the standardization activities can be found in D4.3. We will continue the collaboration regarding the improving of the ID in WP4. However, the collaboration regarding flow-based monitoring of Android devices, will be continued in WP5 and WP6 to develop a modular and flexible service to collect, store and analyze NetFlow data of running applications on the user's device.

The flow-based monitoring of network traffic produced by mobile devices, which is core to this collaboration, contributed to the monitoring activities of WP5. Relating the location information of a device to its network traffic can be beneficial to many network management and measurement applications, including traffic profiling, anomaly detection and provider-independent measurements. The developed approach allows us to better understand Android apps regarding their network flows and usage and it proves to be promising for the automated configuration of mobile networks that concerns WP6. When Metering Processes are running on devices with a (frequently) changing physical location, data analysis applications may need to be aware of these movements since they are likely to affect the behavior of the network in terms of routing, throughput, etc. Thus, configuration policies and actions could be applied to adapt the network according to the observed locations and maintain an acceptable quality of service of running applications on the user's devices. For example, knowing the location of a device at a moment of certain network activities could be used to dynamically reroute its traffic to closer data sources.

### 3.2.3 Mobile Measurements (UZH-JUB-UniBwM-M2)

Quality-of-service (QoS) metrics have been traditionally used to evaluate the perceived quality of services delivered by mobile network operators. However, this metric is not suitable for evaluating the experience of an end-user. Experience of a user is quantified based upon activities such as speed of web page loading, quality of video streaming, or voice quality of Internet-telephony. Due to the temporal and geographical nature of mobile networks, the perceived experience of a user may change based on location and time. Mobile operators may prioritize certain services over others, leading to a service type dependent quality of experience (QoE).

This collaboration aims to develop mechanisms for evaluating activity and protocol based QoE. The aim is to have a method for obtaining a service specific QoE based on active measurements performed in mobile networks. Obtained QoE values will be mapped to mean opinion scores (MOS) and presented on a global map. This not only aids operators in identifying their users' QoE in specific locations, but can also assist users in identifying areas where they might have coverage issues.

---

<sup>6</sup><http://nfsen.sourceforge.net>

<sup>7</sup><http://surfmap.sourceforge.net>

An approach to obtain QoE MOS values based on statistics (bandwidth, latency, signal strength, etc.) has been outlined [19].

This approach has led to the development of the BonaFide+ [20] application. With the initial development of the measurement application complete, measurement data is now being collected and approaches being defined to further refine the QoE calculation approach. Once enough data is collected, the results were analyzed and published. In the meantime, a larger measurement endpoint infrastructure is also being sought after. Assistance of the M-Lab project [21] is likely to be sought, while current endpoints are deployed on EmanicsLab [22] already. A website with information regarding the application, access to source code and collected data is currently being worked upon as well [23].

The aspect of an active collection of network metrics for service quality monitoring and the identification of possible traffic shaping in the network is part of WP5.

### 3.2.4 Schengen Routing (UT-UniBwM-Scheng)

The affair involving Edward Snowden and the National Security Agency (NSA) in 2013 demonstrated that wiretapping large amounts of Internet traffic data was not only possible, but also applied on a regular basis by various intelligence agencies in violation of privacy laws [24]. However, the controversy only came into broader political debate by the time it was alleged that several European state heads had become victims of the wiretapping activities themselves [25].

In the context of the political and technical debate that followed, the idea of Schengen Routing demonstrated to be a possible amendment to protect communications across Europe. The term Schengen refers to the treaty targeted at reducing border controls and implementing a harmonized legal framework [26]. Those countries, who signed the Schengen Treaty, form the Schengen Area. It is important to highlight that the Schengen Area is not equivalent to the European Union (EU), since some countries belonging to the EU are not part of Schengen (e.g., United Kingdom), while Schengen also comprises non-EU countries (e.g., Switzerland).

Schengen Routing refers to the practice of routing Internet traffic between hosts located in the Schengen Area, not leaving the borders of countries part of the Schengen Treaty. Such Internet traffic not leaving the Schengen Area is more difficult to be wiretapped by non-Schengen intelligence agencies, since the Internet traffic remains still unencrypted. However, this traffic remains still vulnerable to wiretapping activities that may occur within Schengen [27]. An implementation of Schengen routing requires the reconfiguration of routing tables and the renegotiation of transit and peering agreements. The effort required depends significantly on the degree to which current routing already complies with Schengen routing or not.

The FLAMINGO work [3] is the first work which measured a Schengen routing compliance through active measurements by analyzing TCP, ICMP, and UDP traffic. It answered the following question: What is the Schengen routing compliance or non-compliance percentage of current traffic among Schengen countries based on the observation of active measurements?. For that a large number of traceroute measurements was executed by applying RIPE Atlas [28] probes located in Autonomous Systems (AS) within Schengen to a well-known host in Switzerland, being part of the Schengen Area. ASes were chosen as the unit of analysis, because ASes are collections of network devices managed by a single administrative authority that can decide to cooperate with government agencies or not. IP addresses of nodes along a network path can be determined by using the traceroute tool. By means of a database, such as GeoLite [29], IP addresses obtained can be related to ASes and countries and, thus, placed in- or outside Schengen. Next to these measurements, a tool termed chkroute has been developed, allowing end-users to find out whether specific routes are Schengen-compliant.

[3] presented key results of a larger-scale measurement conducted to determine the extent to which current routing is Schengen-compliant in Schengen countries. Based on 3388 TCP, UDP, and ICMP traceroute measurements run from RIPE Atlas probes located in over 1100 ASes in the Schengen Area it was found that compliance levels vary substantially among countries and range from 0% (TCP), 0% (UDP) and 0% (ICMP) in the case of Malta to 80% (TCP), 75% (UDP), and 80% (ICMP) in the case of Liechtenstein. The overall compliance levels range from 34.5% (TCP) to 37.4% (UDP) and 39.7% (ICMP). Based on these measurements performed, [3] concludes that Schengen Routing compliance is not achieved in any of the Schengen countries, contradicting the claim that Schengen routing already was a factual reality today, as it has been stated by the Association of the German Internet Industry [30]. Therefore, intelligence agencies still can perform potential wiretapping activities outside the Schengen jurisdiction on traffic originating within and destined to the Schengen Area. [3] and chkroute especially with the data set collected only analyze traffic in the forward direction.

The reverse path may not necessarily be the same [31]. Hence, future work will address the reverse path, too. Furthermore, as only routes originating in Schengen countries targeted at a single node in Switzerland have been analyzed, results may differ, if a target node in another country was chosen. In particular, routes originating in a Schengen country A may be more or less likely than those originating in another Schengen country B to traverse a non-Schengen country depending on the location of the target node. Finally, an additional extension to [3] is to analyze traffic for individual countries in more detail as well as to provide details with respect to countries that cause routes to be non-compliant with Schengen routing. It is also essential to identify those Schengen countries that are exit and entry points for traffic out of the Schengen Area and to examine the reason why this is the case.

### 3.2.5 Machine Learning and Botnet Detection (UniBwM-UT-MLB)

In this joint PhD research activity, the Universität der Bundeswehr München and the University of Twente aim to investigate the Botnet-Phenomenon to improve current network security and defense solutions. Botnets are compromised machines that are remotely controllable by cyber-criminals via a so-called command and control channel (C&C). Botnets provide the basis for many network-based attacks, cyber-criminal activity and the currently growing Crimeware-as-a-Service (CaaS) phenomenon, including for example DDoS attacks, banking-fraud, information theft and extortion.

Current botnet detection approaches suffer from mainly three problems a) frequent updates of the botnet malware causing changed behavior and signatures, b) encryption and other techniques to hide the C&C traffic and c) the global dimension, since bots are spread across country (and network) borders.

This joint research approach aims to address these challenges by using globally distributed network-flow sensors and domain name registration data of the most often used Domains world-wide in combination with self-adaptive detection approaches, which are based on machine learning and artificial intelligence. Consequently, this research collaboration, within FLAMINGO, mainly contributes to work packaged 5 and 6. A more detailed description of the relations to each of the two work packages is given below.

**Relations to WP 5:** The research will make use of distributed flow monitoring and measurement data. Such data will be collected and shared among the flamingo partners. Multiple NRENs and Domain registrars provide data that support this research. In addition, efficient flow querying

solutions and distributed monitoring and measurement facilities, developed within WP5, provide valuable input for this research collaboration.

**Relations to WP 6:** Within this collaboration novel botnet analytic solutions are developed that might enable early detection and network defense systems. Solutions invented within this collaboration can provide input for SDN based network reconfiguration to proactively stop attacks at the borders of a network. Such approaches will benefit from the fact that every remotely coordinated botnet-based attack (e.g. DDoS) needs to be signaled among all the bots in the network. As a consequence, there is usually a time delay between the start of an attack and the attack reaching its apex. This time can be used for example to acquire additional resources for load balancing, changing firewall rules or to deploy adjusted routing rules within a network.

## 4 Lessons Learned from using the RIPE Atlas Platform for Measurement Research

In this section and in the following Sections 5-9, we summarized selected highlights of the research conducted in WP5 during Y3.

RIPE Atlas [32] has deployed around 12.8K dedicated hardware probes and around 109 anchors (as of Feb 2015) all around the globe<sup>8</sup>. Probes perform active measurements to ascertain network connectivity and reachability of the global Internet, while anchors are dedicated servers that can act as sources and sinks of measurement traffic. RIPE Atlas periodically schedules measurements using a batch of several hundred probes against anchors to measure region-based connectivity and reachability. A majority of these probes are running measurements either from the core or from within access networks. A discernible number of probes are also hosted by volunteers within their home networks. All hosted probes are made publicly available for measurement research. These probes in addition to built-in measurements can also run User Defined Measurement (UDM)s. A UDM allows any user registered (around 19K as of Feb 2015) on RIPE Atlas to provision measurements supported by the platform on probes with tailor-made measurement parameters. A registered user spends credits by provisioning a UDM on probes. Credits can be gathered by either hosting a probe (for no purchase cost) or an anchor (for a purchase cost). RIPE Atlas also released (on Feb 2013) a public API that allows one to programmatically provision UDMs. Using these public APIs and credits gathered by hosting probes for multiple years, we were able to provision UDMs on a large sample of probes.

This paper reflects upon the authors experience in using the RIPE Atlas platform for measurement-based research. The authors show how in addition to credits, control checks using rate limits are in place to ensure that the platform does not get overloaded with measurements. They show how the Autonomous System (AS)-based distribution of RIPE Atlas probes is heavily skewed which limits possibilities of measurements sourced from a specific origin-AS. They discuss the significance of probe calibration and how they leverage it to identify load issues in older hardware versions (38.6% overall as of Sep 2014) of probes. They show how performance measurement platforms (such as RIPE Atlas, SamKnows, BISmark and Dasu) can benefit from each other by demonstrating two example use-cases. They also open discussion on how RIPE Atlas deployment can be made more useful by relaying more probe metadata information back to the scientific community and by strategically deploying probes to reduce the inherent sampling bias embedded in probe-based measurement platforms. This summary highlights the relevant findings with respect to the AS analysis performed. For details, we refer the reader to [33].

### 4.1 Heavy-tailed probe distribution

The geographical distribution of the probes provides a decent high-level overview of the coverage of the platform. Although the network coverage map<sup>9</sup> provides a facility to filter probes by AS Number (ASN), the overall distribution of probes across ASes and density of probes within each AS is not well known. Measurements sourced from a specific AS require high probe density to maintain a representative sample, while measurements destined towards a specific AS require diversity of network origins. As such, we performed an experiment to better understand the AS-based distribution of these probes.

<sup>8</sup>The research reported in this section had been published as an editorial note in ACM/SIGCOMM Computer Communication Review [33] V. Bajpai, S.J. Eravuchira, J. Schönwälder *Lessons Learned from using the RIPE Atlas Platform for Measurement Research* SIGCOMM Comput. Commun. Rev., Vol 45, Issue 3, July 2015

<sup>9</sup>[atlas.ripe.net/results/maps/network-coverage](http://atlas.ripe.net/results/maps/network-coverage)

AS Rank	AS (ASN)	# Probes
01	COMCAST (AS7922)	313
02	PROXAD (AS12322)	242
03	LGI-UPC (AS6830)	233
04	DTAG (AS3320)	190
05	ORANGE (AS3215)	124
06	ZIGGO (AS9143)	83
07	XS4ALL (AS3265)	82
08	BT (AS2856)	76
09	UUNET (AS701)	74
10	VIRGINMEDIA (AS5089)	73

Table 8: Distribution of a subset of connected and non-anchored probes (7672) sorted by AS rank as of Feb 2015.

#### 4.1.1 Clustering probes by ASN

We use the RIPE Atlas probe API<sup>10</sup> to capture a list of connected probes in order to later cluster them by their origin AS. The API, however, does not reveal the ASN for all probes. For instance, some probes (2037, 15.9% of all registered probes as of Feb 2015) did not expose either their public IP or their origin-AS. We grabbed the probe IDs of these probes and provisioned a one-off (measurement that runs only once) traceroute measurement. The measurement was scheduled only on a few probes (43 out of 2037) while the rest were deemed disconnected by the scheduler. We identified the origin AS of these probes, and pruned the rest of the disconnected probes out of the list. We also used the mapping in Fig. 4 (described later in the paper) to rule out anchors (109 as of Feb 2015). Going forward, we use the term probe to refer to the connected and non-anchored subset (7672) of all RIPE Atlas probes (12790).

#### 4.1.2 Ranking ASNs by number of probes

We ranked ASNs by sorting them by the number of deployed probes. Table 8 provides a list of top 10 ASes containing the highest number of probes. For instance, Comcast (AS7922) has 313 (out of 7672) probes which contributes to 4% of all probes. The cumulative probes within top 10 AS ranks contribute to 18 of all probes as of Feb 2015. Fig. 5 shows the long-tail probe distribution sorted by AS ranks. A corresponding CDF of this long-tail, shows how probes deployed within AS ranks  $> 101$  have less than even 10 probes. To bring numbers into perspective, if we were to consider 10+ probes as a representative sample within each AS, the number of probes falling within AS ranks  $\leq 101$  would contribute 44.59% (3421 out of 7672) which is less than half of the entire population of probes.

#### 4.1.3 Clustering ASNs by network type

Using PeeringDB, we further mapped ASes hosting the connected probes (7672 as of Feb 2015) by their network type information. PeeringDB<sup>11</sup> is a database holding peering information of participating networks. Aemen Lodhi et al. in [34] show how the information maintained within this

<sup>10</sup>[atlas.ripe.net/api/v1/probe](https://atlas.ripe.net/api/v1/probe)

<sup>11</sup>[peeringdb.com](https://peeringdb.com)




Figure 5: Distribution of a subset of connected and non-anchored probes (7672) sorted by AS rank as of Feb 2015. ASes are ranked by number of probes.




Figure 6: Evolution of probes by network type as mapped by PeeringDB.

database is reasonably representative of network operator peering and is also up-to-date. Fig. 6 shows the evolution of probes by network type over a year. Few spikes occur in the non-profit network type due to a large fraction of probes (with a series of consecutive probe IDs) coming online for a day (or few days) from within the RIPE NCC network. Not all ASes hosting connected probes could be mapped to a network type due to missing AS information (encompassing 33.5% probes as of Feb 2015) in the PeeringDB database. Nevertheless, this mapping provides an indication on which type of networks hold major portions of connected probes. As such, RIPE Atlas is a potential



platform for performing active measurements from within service provider networks.

#### **4.1.4 Skewed distribution of probes**

The RIPE Atlas platform ostensibly appears to have a large number of deployed (12.8K registered as of Feb 2015) probes. However, it turns out that the number of probes available for a measurement study sourced from a specific origin-AS is small. This is due to the skewed distribution of probes which considerably reduces the density of probes behind each AS. In all fairness, the platform was initially designed to measure connectivity and reachability. As such, there has been an inclination to deploy probes to increase coverage (than density) by biasing distribution in favor of under-served ASNs. As a result, the platform is more suitable for performing measurements targeted to a specific destination as it provides diversity of network origins.

## 5 The Internet of Names: A DNS Big Dataset

Next to IP, DNS is arguably the most important infrastructure on the Internet<sup>12</sup>. DNS is pervasive as almost all networked applications and services rely on DNS to map names to IP addresses. Consequently, measuring what is in the DNS can teach us a lot about the state of the Internet. If performed systematically over time, such measurements allow us to observe the evolution of the Internet.

The applications of measuring DNS over time are myriad. An important area of application is network security. Knowledge of what names an IP address mapped to in the past, for example, can be a valuable tool to track malicious activity. There are also many applications in network research. Knowledge of DNS content over time provides empirical data about operational practices and deployment of new protocols. If, for example, we wanted to answer the question how the use of cloud e-mail providers develops over time, knowledge of the DNS is vital (as who handles mail for a domain is configured in DNS through the MX record type).

### Passive DNS

Development of the only existing large-scale approach to DNS measurements, passive DNS (pDNS) [35], was driven by security benefits. Fig. 7 shows an abstract view of the DNS. pDNS typically collects data on the link between recursive caching name servers (“resolvers”) and authoritative name servers (light grey area). Data from pDNS setups<sup>13</sup> can, e.g., be used to track names associated with IP addresses that exhibit malicious activity. From a network research perspective, pDNS is also of interest but it suffers from one problem: it does not provide reliable data over time. This is because 1) pDNS will only record data for domains in which clients behind the resolvers where pDNS data is collected are interested and 2) pDNS has no influence over temporal spacing of queries.

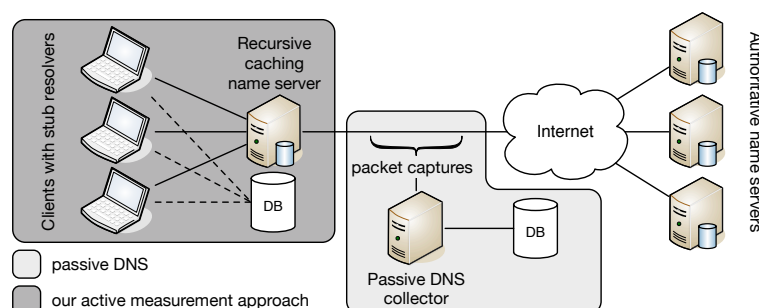


Figure 7: pDNS compared to our approach

### Our approach

Driven by a research need for reliable data from the DNS over time, we have developed a complementary approach to pDNS, based on active measurements. Given the DNS zone files from top-level domains (TLDs) as input, we send a fixed selection of queries for each domain in a TLD once per 24 hours. Effectively, if we compare our approach to pDNS, we control the behaviour of the clients performing queries (shown in dark grey in Fig. 7). It is highly challenging to make such an approach scale. For example, the .com domain alone (the largest TLD on the Internet), already contains >116M names. In the remainder of this poster abstract, we provide a brief outline of our approach and we highlight the potential of the resulting dataset with a case study.

<sup>12</sup>The research reported in this section had been published as a short paper in SIGCOMM 2015 [4] R. van Rijswijk-Deij, M. Jonker, A. Sperotto, A. Pras *The Internet of Names: A DNS Big Dataset* Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication (SIGCOMM 2015) For the associated poster we refer to Appendix A

<sup>13</sup>e.g., DNSDB – <https://www.dnsdb.info/>

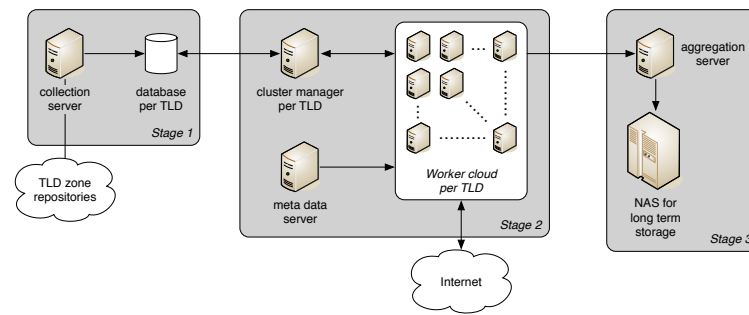


Figure 8: High-level infrastructure overview

TLD	#domains	#workers	avg. time	#queries/day	data/day
.org	10.5 mil.	10	6h45m	127 mil.	2.4GB
.net	15.0 mil.	10	13h30m	181 mil.	3.3GB
.com	116.5 mil.	80	17h30m	1427 mil.	27.2GB
<i>total</i>	142.0 mil.	100	-	1735 mil.	32.9GB

Table 9: Active measurement characteristics

## 5.1 Infrastructure

### 5.1.1 High-level overview

Fig. 8 gives a bird's eye view of our measurement setup. We divide the measurement process into three stages:

- **Stage 1:** input collection – in this stage, we collect the DNS zones for the TLDs to measure. We compute daily deltas and track both the active zone content as well as changes in the zone over time in a database.
- **Stage 2:** main measurement – this is the active measurement stage; we will explain this stage in more detail in Sec. 5.1.2 below.
- **Stage 3:** aggregate and prepare for analysis – in this final stage we convert the output from the measurement to the Parquet columnar storage format, which is well-suited for processing on a Hadoop cluster.

### 5.1.2 Main measurement

Our main active measurement runs on a cloud-based cluster. Every TLD measurement is orchestrated by a cluster management host. This host is responsible for distributing chunks of work, of 100k domains each, to a set of worker nodes. Each worker node runs custom-built software that performs a pre-defined selection of DNS queries for each domain in a chunk of work. Queries are performed against a local DNS resolver instance running on the worker node. Data collected by workers is sent to a central aggregation point for further processing and analysis. Tab. 9 shows an overview of our current setup. It shows the TLDs we measure, the number of worker nodes, the average time to complete a full measurement, the average number of queries per day, and the amount of data collected per day.




Figure 9: Use of cloud mail platforms over time

### 5.1.3 Analysis

As mentioned in Sec. 5.1.1, we store data such that it is well-suited for processing on a Hadoop cluster. We use such a cluster for analysis, defining map/reduce operations on the data, but also more advanced forms of analysis using, e.g., the Impala massively parallel query engine<sup>14</sup>.

## 5.2 Case study

As case study, we have analysed the use of cloud mail platforms over a 50-day period. We focused on two particular platforms, Microsoft's Office 365 and Google Apps. Fig. 9 shows the growth in the fraction of domains per TLD that use either of these platforms. Growth is presented as a percentage relative to the start of the 50-day period. To perform the analysis, our platform processed over 84 billion query results. The full analysis was performed by a single 40-core node in about 7.5 hours. This can easily be improved by running the analysis on a larger Hadoop cluster.

This simple example showcases what can be achieved using our measurement platform and data. The growth in use of, e.g., cloud mail platforms illustrates how the Internet is evolving from every organisation managing its own services to a few large providers offering these services in bulk.

## 5.3 Conclusions and Future Work

We created a unique active measurement infrastructure for the DNS. Our infrastructure actively measures over 50% of the total DNS name space on a daily basis. The resulting dataset enables reliable DNS-based analysis of the evolution of the Internet for the first time. And not only do we measure on a large scale, we have also carefully designed for optimal analysis of the collected data through the Hadoop toolchain. The simple case study included in this abstract showcases use of our dataset. It answers the simple question about cloud mail platforms that we provided as an example in the introduction.

The goal of this poster is to invite other researchers to collaborate with us, to analyse this unique new dataset. To provide insight into its potential, we plan to create a web portal with daily statistics. Furthermore, we have already started several research projects that investigate Internet phenomena and that rely on measurement data from this platform.

<sup>14</sup><http://www.cloudera.com/content/cloudera/en/products-and-services/cdh/impala.html>

## 6 Booters – An Analysis of DDoS-as-a-Service Attacks

In 2012, the Dutch National Research and Education Network, SURFnet, observed a multitude of Distributed Denial of Service (DDoS) attacks against educational institutions<sup>15</sup>. These attacks were effective enough to cause the online exams of hundreds of students to be cancelled. Surprisingly, these attacks were purchased by students from websites, known as Booters. These sites provide DDoS attacks as a paid service (DDoS-as-a-Service) at costs starting from 1 USD. Since this problem was first identified by SURFnet, Booters have been used repeatedly to perform attacks on schools in SURFnets constituency. Very little is known, however, about the characteristics of Booters, and particularly how their attacks are structured. This is vital information needed to mitigate these attacks.

The goal of this research is to create awareness around Booter attacks. In the study, the authors investigate the characteristics of Booter attacks in terms of the volume of generated traffic as well as the service and networking infrastructure used by Booters. Finally, based on measurements, they discuss possible defense mechanisms and the relationship between Booters and DDoS protection services. The authors performed measurements to analyze the attacks generated by Booters on their own infrastructure at the UT network, leading to more than 250 GB of traffic.

This summary highlights some of the relevant analysis results. For the complete analysis, we refer the reader to [36].

### 6.1 Measurement methodology

To investigate the characteristics of Booter attacks, the authors purchased DDoS attacks from 14 Booters that were online and operational on 14 and 15 August 2013. The goal of our experiment was to determine how much traffic Booters are able to generate and the geographical distribution of the systems misused by Booters to perform attacks. In collaboration with SURFnet and the University of Twente (UT), they launched a series of attacks on network infrastructure specifically dedicated to this experiment at the UT. Although our list of Booters at that moment was composed of 21 online Booters, 7 of them had a faulty payment system that did not allow to purchase packages of attacks.

For each of the 14 Booters investigated the authors: 1) create an account, 2) purchase an attack package; and 3) launch UDP- based DDoS attacks against a null-routed IP address at the UT. Although Booters offer several types of DDoS attacks, for these experiments the authors concentrated on volumetric attacks based on UDP because no service running on the target system is required, and the only potential bottleneck is the network link capacity.

During the attacks, we captured raw packet data at the UT using dedicated hardware, capable of capturing traffic at 10 Gbps. To ensure that attacks did not hinder the functioning of UT's or other networks, and that the attack traffic rate remains below the maximum network link rate (10 Gbps), SURFnet and the Computer Security Incident Response Teams (CSIRTs) from SURFnet and the UT were informed and actively collaborated in monitoring the attack traffic.

---

<sup>15</sup>The research reported in this section had been published in IM 2015 [36] J.J. Santanna, R. van Rijkswijk-Deij, R. Hofstede, A. Sperotto, M. Wierbosch, L. Zambenedetti Granville, A. Pras *Booters – An Analysis of DDoS-as-a-Service Attacks*, Proceedings of the IFIP/IEEE International Symposium on Integrated Network Management (IM 2015)

### 6.1.1 Booter attacks

Of the 14 Booters from which we purchased attacks, 5 Booters did not perform the UDP-based attacks that we ordered: 3 of those did not send any traffic, and 2 surprisingly generated a handful of TCP packets. The 9 remaining Booters performed as requested, however, and generated more than 250 GB of traffic.

Although there are several types of UDP-based attacks (*e.g.* amplification attacks, based on NTP, SNMP, DNS, and Echo), our measurements only show 7 DNS-based attacks and 2 attacks involving the CharGen protocol. This observation is in line with current trends described in [37] that show DNS and CharGen as two of the most common types of UDP-based attacks.

Both types of attacks (DNS and CharGen) belong to the class of reflection and amplification attacks. These attacks are based on the principle that an attacker sends a relatively small request to a server, crafted with the spoofed IP address of the intended target (reflection), and for which the response is much larger than the request (amplification). For example, in case of a DNS-based attack, an attacker may send a relatively small DNS query (in the order of 40 – 60 bytes), which may be answered with a large response that can be 4 KB or more in case EDNS0<sup>16</sup> is used. In case of CharGen [39], RFC 864 defines that requests to servers should be answered with a randomly-sized reply up to 512 bytes in size.




Figure 10: Traffic rate of DNS-based and CharGen-based attacks.

Fig. 10 shows the traffic rate generated by the attacks for both DNS-based and CharGen-based attacks. The traffic has been compensated for possible packet drops that occurred due to the collection infrastructure being overloaded. For the compensation algorithm applied to the traffic, we refer the reader to [36].

Surprisingly, we notice a large discrepancy between the maximum allowed packet size as per the CharGen protocol specification in RFC 864 [39] (512 bytes) and what we measure. As shown in Figure 11, for both Booters B8 and B9, the size of packets is randomly distributed in the range of [0, 6956] bytes. Therefore, we suspect that the systems involved in the attacks were running a non-RFC-compliant implementation of the CharGen protocol. To verify this, we first examined the misused systems using `nmap`<sup>17</sup> and we observed that the majority of these systems were running Microsoft Windows.

To verify whether the observed CharGen implementation is specific to MS Windows systems, we installed several recent versions of Microsoft Windows, as well as the reference implementation of

<sup>16</sup>The Extension mechanisms for DNS (EDNS0) [38] allow for - among other things - larger DNS responses (than the originally specified 512 bytes), with the most common maximum size configured set to 4 KB.

<sup>17</sup><http://nmap.org>




Figure 11: Packet size distribution (CharGen).

the `xinetd`<sup>18</sup> daemon on Linux (which includes CharGen). When we tested the protocol in our lab environment, our results confirmed those of the live attacks for the implementations on systems running Microsoft Windows. The maximum CharGen packet sizes measured in our lab environment are remarkable: all Microsoft Windows versions from XP up return messages with a random size of  $[0, 6956]$  bytes. This confirms that the Windows implementations are non-RFC-compliant. In addition, since CharGen is installed as part of the Simple TCP/IP Services, Windows systems may therefore become a powerful base for this type of amplification attack if these services are enabled. Our tests also show that the `xinetd` implementation of CharGen is non-RFC-compliant, although in this case the maximum obtained response size is limited to 1024 bytes, and therefore – on average – 3.4 times smaller than for Microsoft Windows.

<sup>18</sup><http://www.xinetd.org>

## 7 Making the Case for Elliptic Curves in DNSSEC

The Domain Name System Security Extensions (DNSSEC) add authenticity and integrity to the DNS, improving its security <sup>19</sup>.

While DNSSEC can improve the security of the Internet, uptake is still lacklustre. Less than 3% of domains worldwide deploy DNSSEC<sup>20</sup> and at best 13% of clients are protected by DNSSEC validation<sup>21</sup>. We argue that this is partly due to problems with DNSSEC as a technology. Three problems stand out. First, DNSSEC responses are larger and suffer more from IP fragmentation, which impacts availability [41]. Second, DNSSEC's larger responses can be abused for potent denial-of-service attacks [42]. Third, key management in DNSSEC is often complex, which may lead to mistakes that make domains unreachable. These issues raise the question if the benefits of DNSSEC outweigh the disadvantages.

This research argues that the choice for RSA as default cryptosystem in DNSSEC is a major factor in these three problems. Alternative cryptosystems, based on elliptic curve cryptography (ECDSA and EdDSA), exist but are rarely used in DNSSEC. The authors show that these are highly attractive for use in DNSSEC, although they also have disadvantages. To address these, the authors have initiated research that aims to investigate the viability of deploying ECC at a large scale in DNSSEC.

This research highlight reports a sample of the relevant findings in the paper. For the full research, we refer the reader to [40].

### 7.1 Deployment Scenarios

This research considers the deployment scenarios in Table 10. The analysis covers the following ECC implementation choices: a) ECDSA versus EdDSA, b) which curve is used and c) the 'traditional' KSK/ZSK (Key Signing Key/Zone Signing Key) split versus a single Combined Signing Key (CSK). Tab. 10 provides a set of scenarios covering these choices. The rows in Tab. 10 show the implementation choices, the columns provide convenient short names for the scenarios. The scenarios are sorted from most conservative (in terms of existing standards and practices, and with respect to security and proven cryptography) to most beneficial in terms of tackling the issues we identified (but requiring implementation changes or standardisation and relying on more novel cryptographic algorithms). We will test these scenarios using measurements.

<i>implementation choice</i>	<i>ecdsa384</i>	<i>ecdsa256</i>	<i>ecdsa384csk</i>	<i>ecdsa256csk</i>	<i>eddsasplit</i>	<i>eddsacsk</i>
ECDSA vs. EdDSA	ECDSA	ECDSA	ECDSA	ECDSA	EdDSA	EdDSA
Curve	P-384	P-256	P-384	P-256	Ed25519	Ed25519
KSK/ZSK vs. CSK	KSK/ZSK	KSK/ZSK	CSK	CSK	KSK/ZSK	CSK
	<i>most conservative</i>		$\longleftrightarrow$		<i>most beneficial</i>	

Table 10: Deployment scenarios for ECC in DNSSEC

<sup>19</sup>The research reported in this section had been published in ACM/SIGCOMM Computer Communication Review [40] R. van Rijswijk-Deij, A. Sperotto, A. Pras *Making the Case for Elliptic Curves in DNSSEC* accepted for publication in the October 2015 issue

<sup>20</sup><http://www.isoc.org/deploy360/dnssec/statistics/>

<sup>21</sup><http://stats.labs.apnic.net/dnssec/XA>



## 7.2 Fragmentation

### 7.2.1 The problem

DNSSEC responses are larger than regular DNS responses since they include digital signatures. Sometimes this leads to packet fragmentation. Moreover, some DNSSEC-specific query types are particularly at risk from this. The *DNSKEY* query type – crucial for DNSSEC – can have large responses as it includes all public keys used in signing the zone. The authors showed earlier that fragmentation is a big problem [41]. Up to 10% of hosts may be unable to handle fragmented responses. Domains with fragmented DNS responses are effectively unreachable for these hosts. Based on this research, SURFnet<sup>22</sup> enabled “minimal responses”. This tells the name server to respond with the smallest possible answer, preventing most fragmentation. Measurements show this decreases the average response size by 80%. But since it does not set a hard limit on response size, fragmentation can still occur. To quantify fragmentation under “minimal responses”, we examined traffic captured from August 2013 to April 2015 at one of SURFnet’s authoritative name servers. Out of 12 million DNSSEC responses per day, 0.5% suffer fragmentation, for  $\pm 15,000$  distinct query names per day. Fragmentation occurs for all common query types and response statuses (success, referral and non-existence). This underlines that fragmentation remains a problem when deploying DNSSEC.

### 7.2.2 Revisiting Fragmentation using Elliptic Curves

To show the impact of the scenarios in Tab. 10 on fragmentation, we performed two measurements. First, we re-issued queries that resulted in fragmentation in our measurement. We examined if answers to these queries would be fragmented under each of the scenarios and find that even the most conservative scenario (*ecdsa384*) vastly reduces the occurrence of fragmentation. Only 0.3% of previously fragmented responses would still be fragmented under this scenario. Under the most beneficial scenario (*eddsacsk*), less than 0.003% of responses would still be fragmented. To all intents and purposes this is a negligible number.




Figure 12: CDF for DNSKEY response sizes

The second measurement examined the effect of our scenarios on DNSSEC-specific query types that earlier research [41, 42] shows suffer from fragmentation. Particularly the response to a *DNSKEY* query may suffer from fragmentation. We examined *DNSKEY* responses for the 0.5 million .com, .net and .org domains with DNSSEC and calculated the response sizes under our scenarios. Fig. 12 shows the top 10% of a CDF plot of the results. The figure shows that 6.5% of current

<sup>22</sup>The National Research and Education Network in the Netherlands (<http://www.surfnet.nl/en/>).

DNSKEY responses exceed the IPv6 minimum MTU and that 0.6% exceed the MTU of Ethernet. It also shows that even switching to the most conservative scenario (ecdsa384) effectively stops fragmentation. But even more remarkable is that two CSK scenarios (ecdsa256csk, eddsadsk) are so effective that the majority of DNSKEY responses would fit in a classic DNS datagram of 512 bytes. We briefly examined the long tail that exceeds this classic DNS limit for these two scenarios and found that simple configuration changes – e.g. enabling “minimal responses” – can make all answers fit in a classic DNS datagram under these two scenarios.

## 8 Mitigation of Topological Inconsistency Attacks in RPL based Low Power Lossy Networks

### 8.1 Introduction

#### 8.1.1 Context

The Routing Protocol for Low-power Lossy Networks (RPL) [43], designed for constrained devices and networks, is expected to find application in multiple areas of the Internet of Things (IoT). Being suitable for various fields like, Industrial Networks [44], Home and Building Automation [45] and Advanced Metering Infrastructure (AMI) Networks [46], it is evident that RPL will be exposed to multiple different operating scenarios, some of which will expose it to malicious attacks.

#### 8.1.2 The RPL protocol

The Routing Protocol for Low-power Lossy Networks (RPL) has been designed by the IETF [43] to address resource constraints of embedded devices. This protocol enables a distance-vector routing based on IPv6. RPL forms a loop-free tree like topology termed a Destination Oriented Directed Acyclic Graph (DODAG). A network can operate one or more RPL instances which consist of multiple DODAG graphs. When a loop occurs, RPL provides the *data path validation* mechanism to detect and repair rank related DODAG inconsistencies. This mechanism works by carrying the following flags in the RPL IPv6 header options [47] of multi-hop data packets:

- The '*O*' flag — indicates the expected direction of a packet. When set, the packet is intended for a descendant. Otherwise it is intended for a parent, towards the DODAG root.
- The '*R*' flag — indicates that a rank error was detected by a node forwarding the packet. A mismatch between the direction indicated by the '*O*' flag and the rank of sending/forwarding node causes the flag to be set.

A DODAG inconsistency exists if the direction indicated by the '*O*' flag does not match the rank relationship of the node from which the packet was received [43]. The '*R*' flag is used to repair this problem by setting it, in case it was not set previously, and forwarding the packet. Upon receiving a packet with the '*R*' flag already set an inconsistency is detected, the packet is discarded and the trickle timer used by RPL is reset [48]. This detection mechanism can be exploited by a malicious node to attack the network.

#### 8.1.3 Attack description

The data path validation can be misuse either to harm a targeted node directly, or to manipulate packet headers and cause the next-hop node to drop the modified packet.

A malicious intruder can directly attack its neighborhood by sending packets that have the '*R*' flag and the wrong direction set. For instance, if a parent is targeted, the attacker can send packets with the '*O*' and '*R*' flags set, since packets with '*O*' flag are intended for descendant nodes. The parent will detect an inconsistency and thus, drop the packet and restart the trickle timer. This causes control messages to be sent more frequently which leads to local instability in the network. This increased control message overhead reduces channel availability and increases energy consumption which can lead to a shortened network lifetime in case nodes are battery operated. Since

nodes in RPL networks are likely to be resource constrained, they are unlikely to support multi-tasking or large packet buffers. As such, time spent on processing malicious packets could lead to loss of genuine ones.

A malicious intruder can also modify the IPv6 header of packets it forwards such that the 'R' flag and the 'O' flag representing the wrong direction are set. The receiving node assumes that a DODAG inconsistency has taken place and discards the packet. As a result, the malicious node succeeds in forming a black-hole at the next-hop node. This attack could either be carried out on all packets forwarded by the malicious node, or selectively based on source, destination, or even type of message. In general this approach is a good strategy for the attacker to force another node to drop the packets. Furthermore, if the control packets originating from the malicious node are normal, then the malicious activity is completely hidden. In this scenario, not only does the delivery ratio decrease, but the control overhead of RPL nodes also increases along with deteriorating channel availability and increasing energy consumption.

## 8.2 Results on attack mitigation and detection

### 8.2.1 Mitigation approaches

Three mitigation approaches are presented in [49]. The first solution called default DODAG inconsistency attack mitigation is based on a fixed threshold set to 20 as proposed in [47]. A local counter keeps tracks of number of 'R' flag packets received. Upon reaching the threshold, malformed packets are dropped but the trickle timer is not reset. This approach limits the impact of a DODAG inconsistency attack, but the value of the threshold is arbitrarily set. However it does not mitigate the indirect scenario attack. In order to take into account the current network state and react to varying attack patterns we developed an adaptive threshold (AT) [50], which determines when to stop resetting the trickle timer. Instead of a constant, a decreasing exponential function is used with fixed parameters. The adaptive threshold causes the threshold to change based on network conditions. If an attacker is aggressive, the threshold drops quickly and increases slowly once the attacks stop. To counter the packet manipulation DODAG inconsistency attack, an extension was made to the adaptive threshold. Nodes behave normally until the number of messages indicating an inconsistency becomes greater than the threshold obtained from the function. This situation indicates either an attack against the node, or malfunction of the node forwarding such packets. To rectify the situation, the node clears the 'O' and 'R' flags before forwarding the packets normally. We have improved our adaptive threshold mitigation approach via the design of a fully dynamic threshold, which is based on network characteristics [49]. This dynamic threshold (DT) is similar to the previous one (decreasing exponential function), however, the parameters for this function are based on node specific characteristics (number of neighbors). The algorithm used to detect the attack has been adapted too. Multiple packets with an 'R' flag can arrive as a result of the same inconsistency. As such resetting the trickle timer each time an 'R' flag packet is received leads to unnecessary overhead. To avoid this situation, a convergence timer is introduced, it is used to ensure that no further trickle timer resets take place within the amount of time it takes for an RPL neighborhood to typically converge. A new counter that keeps track of the number of trickle timer resets is introduced and compared with the calculated threshold to determine when the trickle timer should be reset. The solution was also adapted to counter the packet manipulation scenario by allowing a node to forward packets with the inappropriate flags under certain conditions.




Figure 13: Time-line of outgoing control messages overhead experienced by the attacked node, when the rate of attack is varied from 20 to 1800 packets per hour. Comparison of (a) fixed threshold and adaptive threshold time-lines, (b) fixed threshold and dynamic threshold.

### 8.2.2 Results

In a first series of experiments we have studied performance of the different mitigation approaches in a direct attack scenario. Figure 13 compares the fixed threshold approach to adaptive and dynamic thresholds approaches. Both solutions outperforms the fixed threshold by reducing the control message overhead between 20%-50%. In a second series of experiments we have evaluated our approaches in the packet manipulation scenario. In case of black-hole attack scenarios, which are not mitigated by the default RPL approach, our method can improve the delivery ratio to 99% as against 33% for the default RPL mitigation approach. The performance of our two approaches is quite similar in case of aggressive attacks, however, in all other scenarios the dynamic threshold outperforms the adaptive, thereby making it more suitable for use. Moreover, the dynamic threshold does not require any empirically learned values to be configured.

We have also showed that our solutions have a limited cost for the nodes and can achieve significant energy savings.

## 9 Impact of Packet Sampling on Link Dimensioning

### 9.1 Research Problem

Link dimensioning is often used by network operators aiming at optimal bandwidth provisioning or for network planning. Commonly, network operators use traffic averages obtained by polling SNMP MIBs (e.g., interface octet counters) every couple of minutes [51], with graphical support provided by tools such as MRTG (*Multi Router Traffic Grapher*). These approaches, often referred to as *rules of thumb*, simply add to the traffic average a fixed amount of bandwidth as a safety margin that may depend on various factors such as time of the day. However, since the traffic averages are updated only every few minutes (e.g., 5 minutes), traffic fluctuations that happen at much shorter timescales are completely overlooked. This can have a significant impact on applications that are sensitive to bandwidth fluctuations at small time scales, such as real-time video streaming. To account for these short-term fluctuations, alternative approaches to the rule of thumb demand traffic measurements at the packet level. The problem is that continuous packet capturing does not scale well with the massive volume of traffic in today's high-speed links, requiring dedicated and expensive measurement equipment. This makes packet-based approaches likely impracticable due to operational and economical limitations.

To cope with the increasing volume of traffic, many network operators deploy traffic monitoring based on *packet sampling*. For example, CERN (*European Organization for Nuclear Research*) [52] and AMS-IX (*Amsterdam Internet Exchange*) [53] use sFlow [54] to monitor and measure the huge volumes of their network traffic. Packet sampling aims at reducing the excessive load of traffic measurement, storage and processing, while still having highly granular data. This raises the question whether the advantages of packet sampling can be leveraged for link dimensioning.

### 9.2 Contributions

In this work the authors demonstrated the feasibility of using sampled data for link dimensioning purposes. In particular, the work investigated the Bernoulli, 1-in- $N$  [55] and sFlow [54] sampling strategies. To do so, the link dimensioning formula proposed in [56, 57] was used. The authors proposed approaches to accurately estimate traffic variance from sampled data, which is used as input to the used dimensioning formula. In addition, the authors also studied the impact of the loss of individual packet timestamps caused by the sFlow exporting process on link dimensioning.

The experiments in this research used real network traffic traces that were collected at multiple locations around the globe. These are packet captures, which allowed for the validation of results against empirically defined ground-truth.

### 9.3 Proposed Approaches

The link dimensioning formula used in this work aims at “link transparency” by guaranteeing that users – almost – never perceive network performance degradations due to lack of bandwidth. To statistically ensure transparency to users, the provided link capacity  $C$  should satisfy  $P\{A(T) \geq CT\} \leq \varepsilon$ , where  $A(T)$  denotes the total amount of traffic arriving in intervals of length  $T$  and  $\varepsilon$  indicates the probability that the traffic rate  $A(T)/T$  is exceeding  $C$  at timescale  $T$ . The dimensioning formula requires that traffic aggregates  $A(T)$  at timescale  $T$  are *Normal distributed* and *stationary*. The link capacity  $C(T, \varepsilon)$  needed to satisfy the transparency condition can be calculated by

$$C(T, \varepsilon) = \rho + \frac{1}{T} \sqrt{-2 \log(\varepsilon) \cdot v(T)} \quad ,$$

where the mean traffic rate  $\rho$  is added with a “safety margin” that depends on the variance  $v(T)$  of  $A(T)$ . Notice that  $\rho$  and  $v(T)$  can be easily calculated from non-sampled packet captures by, respectively,

$$\rho = \frac{1}{nT} \sum_{i=1}^n A_i(T) \quad \text{and} \quad v(T) = \frac{1}{n-1} \sum_{i=1}^n (A_i(T) - \rho T)^2, \quad ,$$

where  $A_i(T)$  is the amount (in bytes) of observed traffic in time interval  $i$  of length  $T$  and  $n$  the number of monitored intervals. From sampled data, only a fraction of the traffic is available to compute the above statistics and, therefore, estimations should be proportionally scaled according to the sampling rate used. Consider that  $r$  is the ratio between the total number of monitored packets and the number of sampled packets (i.e.,  $r = 10$  for 1:10 sampling), and  $\rho'$  be the mean traffic rate of the sampled traffic and let  $A'_i(T)$  be the amount of sampled traffic (in bytes) observed in time interval  $i$  of length  $T$ . The original amount of traffic  $A_i(T)$  in that interval can be estimated by  $A_{i,est}(T) = r \cdot A'_i(T)$ . Hence, the original mean traffic and variance can be estimated by, respectively

$$\rho_{est} = \frac{r}{nT} \sum_{i=1}^n A'_i(T) \quad \text{and} \quad v_{est}(T) = \frac{r^2}{n-1} \sum_{i=1}^n (A'_i(T) - \rho' T)^2 \quad .$$

For the traffic variance, however, the approach to account for sampling is not as straightforward as for estimating mean traffic. To account for the additional variance introduced by the sampling process, the authors developed estimators based on the different sampling algorithms. For the case of Bernoulli sampling, the variance can be estimated by

$$v_{bern}(T) = v_{est}(T) - (r-1)E[P]E[S^2] \quad ,$$

where  $r = 1/p$ ,  $v_{est}(T)$  is the naive scaling estimation described above,  $E[P]$  is the average number of packets per time interval before sampling, and  $E[S^2]$  is the second moment of the packet size before sampling.  $E[P]$  can be estimated by multiplying the measured average number of sampled packets per time interval by  $r$ .  $E[S^2]$  can be estimated directly from the sizes of the sampled packets since the sampled packets have the same size distribution as the non-sampled packets according to our assumptions. Alternatively, traffic models could be used to obtain  $E[S^2]$ .

A mathematical treatment of 1-in- $N$  sampling is much more complex than for Bernoulli sampling. The sampling window of  $N$  packets can stretch over several intervals  $T$ , making the sampled traffic  $A'_i(T)$  and  $A'_{i+1}(T)$  in adjacent time intervals dependent. Similar difficulties arise for sFlow sampling. Hence, we simplify the problem and assume that  $1/N$  of the packets of each interval are sampled. Furthermore, we assume again that the numbers of packets per interval in the original traffic stream are i.i.d. like  $P$  and packet sizes are i.i.d. like  $S$ , and  $P$  and  $S$  are independent. Under these assumption, the number of sampled packets  $P'_i$  in time interval  $i$  is  $P'_i = P_i/r$  with  $r = N$  (we ignore the problem that  $P_i/r$  might not be a natural number). The estimation of variance from data sampled 1-in- $N$  can be calculated by

$$v_N(T) = v_{est}(T) - (r-1)E[P]Var[S] \quad ,$$

where  $E[P]$  is the average number of packets per time interval before sampling, and  $Var[S]$  is the variance of the packet size before sampling. Again,  $E[P]$  can be estimated from the number of sampled packets and  $E[S^2]$ , and consequently  $Var[S]$ , can be estimated directly from the sizes of the sampled packets or, alternatively, from traffic models. The estimation of variance from sampled data obtained from sFlow is done using the same estimator formula as for 1-in- $N$  sampling. For more details on the mathematical models, please refer to the published paper [58].

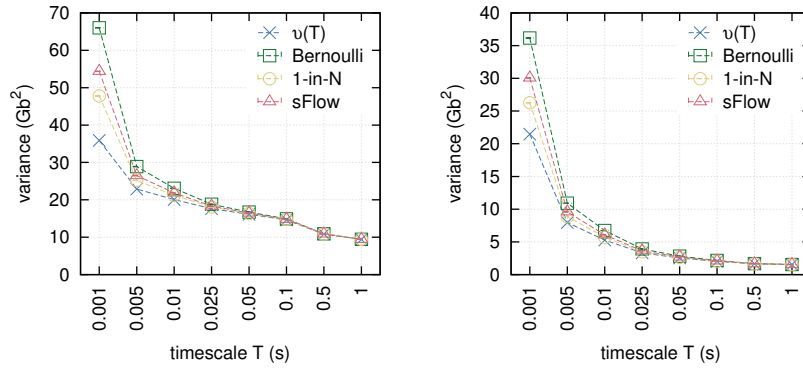


Figure 14: Difference between the average traffic variance calculated from 10 runs of sampling for each algorithm. Example using randomly chosen traces and sampling 1:10.

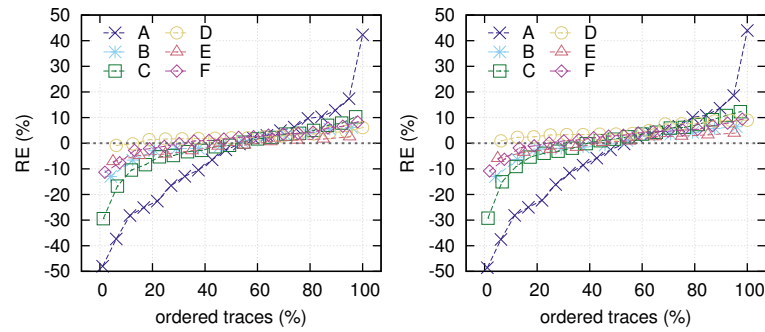


Figure 15: Relative error of estimations as compared to ground-truth for all traces per location. Traces sampled using 1-in- $N$  algorithm and  $T = 1$ s. Left: sampling 1:10. Right: sampling 1:100.

## 9.4 Results and Validation

Fig. 14 shows the average traffic variance computed from 10 runs of sampling for each technique always using two example traffic traces from different locations. The variances were computed using the proposed estimators. Note that the actual traffic variance  $v(T)$  is also plotted for matters of comparison. Although very small and difficult to see, the standard deviation is plotted as error bars in this figure. For both example traces, at any  $T$  or sampling algorithm, the standard deviation is smaller than 1% of the average variance. This confirms that, even with the random nature of the sampling algorithms, several runs still yield very close results. Concerning the traffic variance, the difference between the three sampling algorithms is larger at  $T = 1$ ms, although still not very significant. At any other  $T$ , this difference is negligible or inexistent.

Since results for the estimation of the variances do not yield significant differences among the different sampling algorithms and their respective estimators, the large-scale validation of the estimators proposed in this work has been done only for traces sampled using the 1-in- $N$  strategy.

To validate the developed estimators against the whole dataset (consisting of traces from 6 different locations, named A to F), the authors compared the estimated required capacity  $C$  for each trace with an empirically defined required capacity (ground-truth). Figure 15 shows the relative error ( $RE$ ) obtained for each trace, i.e., how much the estimated  $C$  differs from the ground-truth. Note that results are normalized by the ground-truth capacity. That is, the desirable  $RE$  is close to zero or slightly higher than zero. In case  $RE < 0$  the estimated  $C$  actually underestimates the required capacity of the trace. If  $RE > 0$  the required capacity is overestimated, which is desirable until certain threshold. That is, in an operational environment, unnecessarily overestimated traffic might



result in eventual waste of link resources. The plots in Figure 15 show that the majority of the traces had an estimated capacity  $C$  within reasonable bounds, which is  $\pm 15\%$  the ground-truth. For example, approximately 74% of all traces are within a boundary of  $RE = \pm 10\%$  for  $T = 1\text{s}$ . The worst cases are few traces from location  $A$  for which  $RE$  was up to  $\pm 50\%$ .

## 9.5 Concluding Remarks

The easy-to-use and accurate approaches proposed by the authors expand the applicability of link dimensioning. The solutions proposed can be used in multiple networking operations, for example, network capacity planning, services of on-demand bandwidth reservation and as support to adaptation/changes in the network topology by auto-configuration loops in self-managed networks.

However, the combination of short timescales with inappropriate sampling rate might result in excessive overestimation. To further mitigate overestimation of required capacity, the assumptions in the proposed mathematical models could be sharpened to better represent actual network traffic. For example, accounting for dependent packet numbers and sizes could potentially result in even more accurate estimations of required capacity. Such research is, however, considered future work.

## 10 Integration of EU research

The overall FLAMINGO activities in the context of the integration of the EU research landscape are reported in D3.4. Some of those activities deal with the topics of data and monitoring, and therefore involve WP5. This section reports on the International activities (Section 10.1) and the collaborations with other EU projects and institutions (Section 10.2) that involved WP5.

### 10.1 International Activities

In collaboration with WP2, WP3 and WP4, the consortium has been actively involved in several international activities focussing on diverse aspects of network and service monitoring. We report here the ones that have been relevant, in topic and for the partners participation, to WP5, and we refer to the respective WP2, WP3 and WP4 deliverables for the general overview of these activities:

- The Dagstuhl seminar **Global Measurements: Practice and Experience**<sup>23</sup> (Schloss Dagstuhl, January 4-7, 2016). The seminar focuses on large-scale Internet measurements and their impact on network operation and services, as well as, on the large amount of data that such measurements infrastructures produce. This Dagstuhl Seminar is a follow-up of Dagstuhl Seminar 13472 Global Measurement Frameworks that took place in November 2013, and it aims at discussing the practical experience gained with global measurement frameworks. The seminar is organized in collaboration with the Leone project<sup>24</sup> by: Arthur W. Berger (Akamai Technologies, US), Philip Eardley (British Telecom R&D, GB), Jörg Ott (Aalto University, FI), Jürgen Schönwälder (Jacobs University Bremen, DE).
- Filip De Turck (iMinds), Marinos Charalambides (UCL), Jérôme François (INRIA), Corinna Schmitt (UZH), Ricardo Schmidt (UT), Tim Wauters (iMinds) have taken part in the organization of the **9th International Conference on Autonomous Infrastructure, Management and Security** (AIMS 2015, Gent, Belgium), as Conference Chair, TPC Co-Chair, PhD Student Workshop co-Chairs, and Lab Co-Chairs, respectively. In addition, we highlight the WP-5 related Lab sessions “Map-Reduce and Hadoop” by J. François (INRIA) and “Powering Monitoring Analytics with ELK Stack” by A. Lahmadi and F. Beck (INRIA). Finally, several members of the consortium have acted as TPC Chairs.
- The **6th Workshop on the Usage of NetFlow/IPFIX in Network Management** has taken place during the IETF 93 meeting (Prague, 19-24 July 2015), as the 37th Network Management Research Group (NMRG) meeting, and was organized by Ramin Sadre (UC Louvain) and Ricardo Schmidt (UT). The workshop investigates how NetFlow/IPFIX is used in practice in various aspects of network monitoring and management and it aims at bringing together researchers, operators and manufacturers to exchange their hands-on experience. The workshop was structured around 9 presentations and had 34 participants. UT, UZH and UniBwM actively took part in the workshop by presenting flow-based research activities.
- On September 7, 2015, hosted at the UT, a combined **SAND/FLAMINGO meeting** has taken place. The Self-managing Anycast Networks for DNS (SAND)<sup>25</sup> project aims at developing solution(s) for supporting the complex management of anycast DNS, and it is collaboration between UT, SIDN labs<sup>4</sup> (the organization managing the .nl DNS zone) and NLnetLabs<sup>26</sup> (a

<sup>23</sup><http://www.dagstuhl.de/16012>

<sup>24</sup><http://www.leone-project.eu/>

<sup>25</sup><http://www.sand-project.nl/>

<sup>26</sup><http://nlnetlabs.nl/>

Dutch organization involved in several DNS-based project and software, e.g. Unbound). The meeting, organized as a set of technical presentations, was motivated by the increased interest in DNS related measurement and research activities in FLAMINGO, and aimed sharing hands-on experience with DNS measurements from both projects.

- UT, UniBwM and UZH have organized, in collaboration with CAIDA, the **IJNM special issue Measure, Detect and Mitigate Challenges and Trends in Network Security**. Goal of the special issue is to bring together contributions in the field of network security that have a strong measurement component. The special issue has attracted 14 submissions and, after a thorough review process, 6 papers have been published, among which one FLAMINGO publication [49].

## 10.2 Collaborations with Other EU Projects and Institutions

Given the focus of WP5 on network and service measurements, FLAMINGO has also collaborated with other EU project that are active on the topic of measurements. In particular, FLAMINGO collaborated with:

- The FP7 Project Leone<sup>24</sup> – From global measurements to local management (grant no. 317647).
- The FP7 Project mPlane<sup>27</sup> – Building an Intelligent Measurement Plane for the Internet (grant no. 318627).
- The FP7 Project SmartenIT<sup>28</sup> – Socially-aware Management of New Overlay Application Traffic combined with Energy Efficiency in the Internet (grant no. 317846).
- The STREP project SALUS<sup>29</sup> – Security and interoperability in next generation PPDR communication infrastructures (grant no. ICT-313296)

FLAMINGO and Leone contributed to the organization of the Dagstuhl seminar **Global Measurements: Practice and Experience** (see Section 10.1). The collaborations with other EU projects also had as an outcome several published and submitted papers, as highlighted in Section 2.1.

---

<sup>27</sup><http://www.ict-mplane.eu/>

<sup>28</sup><http://www.smartinit.eu/>

<sup>29</sup><http://www.sec-salus.eu/>

## 11 Conclusions

This deliverable described WP5 achievements in the field of network and service monitoring for the third year of the projec. The WP has targeted and exceeded the S.M.A.R.T. objectives, which focused on the integration of PhD students and the scientific output. Besides the S.M.A.R.T. objectives, the WP has also demonstrated active research in the WP-specific objectives, among which research topics focusing on security are particularly active.

After the first year comment from the reviewers, Y2 has focused on publishing in high-quality conferences and journals. The fruits of such decisions are fully visible in Y3, in which we focused even more decisively on high-quality venues. Visible results are the *IMC Community Contribution Award* awarded to the paper [1] for the best public dataset. The same paper has also been awarded a *IRTF Applied Networking Research Prize 2015*, and publications such as the ones in ACM IMC 2014, SIGCOMM 2015 (short paper), ACM Computer Communication Review, and TNSM. In addition, several papers are currently under reviews in venues such as INFOCOMM 2016, IEEE Communications Magazine and IEEE Journal on Selected Areas in Communications.

WP5 is also focusing in collaborations. For WP5 publications, 21 have been achieved in collaboration with other EU projects and institutions. In addition, some of the measurements activities, in particular the ones around DNS (the Internet of Names and the collaboration with the SAND project) have open concrete possibilities of collaborations with institutions like Center for Applied Internet Data Analysis (CAIDA, UCSD, USA) and the Information Sciences Institute (ISI, USC, USA). Long term visits have already been planned and possibilities in terms of data sharing are being investigated.

In terms of number of publications, the scientific outcome of the WP, both relatively to the number of publications is outstanding. The project has published 73 papers, of which several have a strong measurement component.

Y3 has also seen the emerging of activities clearly aimed at long term measurements, such as the *Internet of Names* (Section 5) and the still growing *Internet traffic statistics*. Such measurements are of particular interest because they not only will progress in Y4, but they will most likely live on after the end of the project, with a clear impact on the research community as well as society at large.

## Abbreviations

<b>AI</b>	Artificial Intelligence
<b>AMI</b>	Advanced Metering Infrastructure
<b>API</b>	Application Programming Interface
<b>AS</b>	Autonomous System
<b>ASN</b>	Autonomous System Number
<b>CharGen</b>	Character Generator Protocol
<b>CLI</b>	Command Line Interface
<b>DDoS</b>	Distributed Denial of Service attack
<b>DNS</b>	Domain Name System
<b>DNSSEC</b>	Domain Name System Security Extensions
<b>DODAG</b>	Destination Oriented Directed Acyclic Graph
<b>FN</b>	False Negatives
<b>FP</b>	False Positives
<b>ECC</b>	Elliptic Curve Cryptography
<b>ECDSA</b>	Elliptic Curve Digital Signature Algorithm
<b>EdDSA</b>	Edwards-curve Digital Signature Algorithm
<b>EDNS</b>	Extension mechanisms for DNS
<b>EU</b>	European Union
<b>Gbps</b>	Giga bits per second
<b>HAS</b>	HTTP Adaptive Streaming
<b>HTTP</b>	Hyper-text Transfer Protocol
<b>HTTPS</b>	Hyper-text Transfer Protocol Secure
<b>ICMP</b>	Internet Control Message Protocol
<b>ID</b>	Internet Draft
<b>IDS</b>	Intrusion Detection System
<b>IDMEF</b>	Intrusion Detection Message Exchange Format
<b>IETF</b>	Internet Engineering Task Force
<b>IP</b>	Internet Protocol
<b>IPFIX</b>	Internet Protocol Flow Information Export

<b>IRTF</b>	Interent Research Task Force
<b>ISP</b>	Internet Service Provider
<b>ITSA</b>	Internet Traffic Statistics Archive
<b>KSK</b>	Key Signing Key
<b>MIB</b>	Management Information Base
<b>MON</b>	Mean Opinion Score
<b>MRTG</b>	Multi Router Traffic Grapher
<b>MTU</b>	Maximum Transmission Unit
<b>NFQL</b>	Network Flow Query Language
<b>NMRG</b>	Network Management Research Group
<b>NREN</b>	National Research and Education Network
<b>NSA</b>	National Security Agency
<b>NTP</b>	Network Time Protocol
<b>pDNS</b>	passive DNS
<b>QoE</b>	Quality-of-Experience
<b>QoS</b>	Quality-of-Service
<b>RFC</b>	Request for Comment
<b>RPL</b>	Routing Protocol for Low power and Lossy Networks
<b>SIDN</b>	Stichting Internet Domeinregistratie Nederland
<b>SDN</b>	Software-Defined Networking
<b>S.M.A.R.T</b>	Specific Measurable Achievable Relevant Timely
<b>SNMP</b>	Simple Network Management Protocol
<b>SSH</b>	Secure SHell
<b>TCP</b>	Transmission Control Protocol
<b>TLD</b>	Top-Level Domain
<b>TN</b>	True Negatives
<b>TP</b>	True Positives
<b>UDM</b>	User Defined Measurement
<b>UDP</b>	User Datagram Protocol
<b>VoD</b>	Video on Demand
<b>ZSK</b>	Zone Signing Key

## A Internet of Names - Poster

This appendix includes the poster titled “The Internet of Names: a DNS Big Dataset – Actively Measuring 50% of the Entire DNS Name Space, Every Day”, by R. van Rijkswijk-Deij, M. Jonker, A. Sperotto, and A. Pras (UT), which has been presented at SIGCOMM 2015 [4].

The Domain Name System (DNS) is part of the core infrastructure of the Internet. Tracking changes in the DNS over time provides valuable information about the evolution of the Internet’s infrastructure. Until now, only one large-scale approach to perform these kinds of measurements existed, passive DNS (pDNS). While pDNS is useful for applications like tracing security incidents, it does not provide sufficient information to reliably track DNS changes over time. We use a complementary approach based on active measurements, which provides a unique, comprehensive dataset on the evolution of DNS over time. Our high-performance infrastructure performs Internet-scale active measurements, currently querying over 50% of the DNS name space on a daily basis. Our infrastructure is designed from the ground up to enable big data analysis approaches on, e.g., a Hadoop cluster. With this novel approach we aim for a quantum leap in DNS-based measurement and analysis of the Internet.

# The Internet of Names: a DNS Big Dataset

Actively Measuring 50% of the Entire DNS Name Space, Every Day

Roland van Rijswijk-Deij†  
r.m.vanrijswijk@utwente.nl

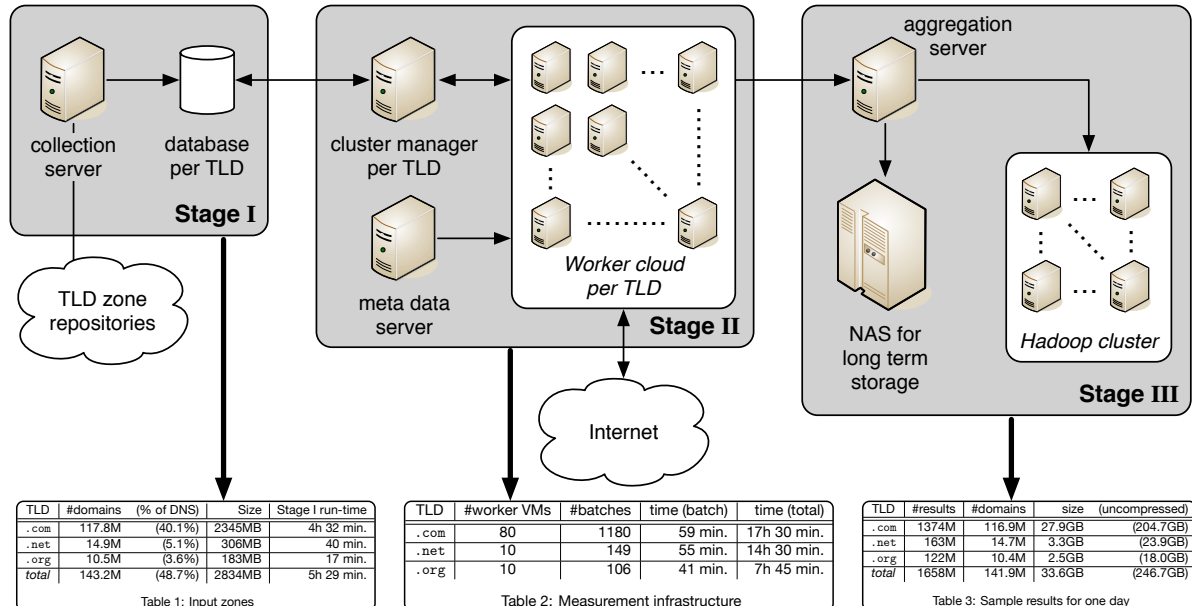
Mattijs Jonker\*  
m.jonker@utwente.nl

Anna Sperotto\*  
a.sperotto@utwente.nl

Aiko Pras\*  
a.pras@utwente.nl

\* Design and Analysis of Communication Systems (DACS),  
Faculty of Electrical Engineering, Mathematics and Computer Science,  
University of Twente, Enschede, The Netherlands

† SURFnet bv, Utrecht, The Netherlands



## Why Measure DNS?

The Domain Name System is used by almost every networked service. It maps human readable names to IP addresses, but also, for instance, which hosts handle e-mail for a domain or information about PKI certificates. Thus, measuring what is in the DNS can tell us a lot about the state of the Internet.

## Our Goals:

We want to measure the DNS to track the evolution of the Internet over time. We therefore want to:

- Measure every domain in the main top-level domains
- Collect data for each domain at least once every 24 hours
- Store at least 1 year of data

## Query Types




- **SOA** DNS zone configuration
- **A** IPv4 address (for apex, 'www' and 'mail')
- **AAAA** IPv6 address (for apex, 'www' and 'mail')
- **NS** Authoritative name servers
- **MX** Mail exchangers
- **TXT** Arbitrary text, used for SPF, DKIM, DMARC, proof of domain ownership, ...
- **DS** Delegation Signer (secure DNSSEC delegation)
- **DNSKEY** Public keys used in DNSSEC signing
- **NSEC(3)** Authenticated Denial-of-Existence (used by DNSSEC)

## Data Sharing

We are working on a public web portal to share aggregate datasets (e.g. number of domains with at least one AAAA record per country, ...). Next to that, we are setting up a programme for researchers to visit our group to use the data we collect (which cannot be made public due to contractual constraints). **If you are interested in visiting us, please contact us via e-mail!**

## Future Work

- Extend the measurement with additional (cc)TLDs
- Collaborate with CSIRT teams to explore security applications of the data



## Example 1: Top Mail Handlers in .com

### Takeaways:

- Hosters/registrars dominant
- Cloud e-mail services clearly visible
- "Domainers" significant presence

## Example 2: Growth of Cloud E-mail

### Takeaways:

- Use of cloud e-mail is growing
- Microsoft Office 365 fastest grower
- Windows Live (Hotmail) dying off

## Example 3: Anomalies for IPv6 in SPF

### Takeaways:

- Dataset well-suited for time series
- Anomalies lead to discoveries
- Sparks new research avenues



## **B Determining the State of Security in the IPv6 Internet**

This appendix include the poster titled “Determining the State of Security in the IPv6 Internet”, by L.Hendriks, A. Sperotto and A. Pras (UT). The poster has been presented at the Terena Networking Conference 2015 [59]. The poster has been presented also at the Traffic Monitoring and Analysis PhD School 2015.

Networks are transitioning from IP version 4 to the new version 6. Fundamental differences in the protocols introduce new security challenges with varying levels of evidence. As enabling IPv6 in an existing network is often already challenging on the functional level, security aspects are overlooked, even those that are emphasized in literature. Assuming the security solutions we know from IPv4 networks are applicable and suffice is unproven and ignorant, possibly leaving IPv6 networks vulnerable to unseen threats. By performing distributed network measurements in production networks, based on both honeypot and attack tool analysis, we determine the current state of security in the IPv6 domain. With the inevitable switch to the new protocol version, assessing the applicability of existing security approaches and determining the requirements for new solutions becomes a necessity.

# Determining the state of Security in the IPv6 Internet

Luuk Hendriks,  
Anna Sperotto,  
Aiko Pras  
*Design and Analysis of  
Communication Systems*  
University of Twente  
luuk.hendriks@utwente.nl

**Security research is often performed and validated based on IPv4-only data.**


Assuming results and approaches apply for its successor IPv6 is a danger in itself. While many works describe IPv6 vulnerabilities in a theoretical way, real large-scale measurements have been limited to a handful of darknets, yielding minimal results, thus providing insufficient insights.

**Distributed measurements characterize the current IPv6 security landscape.**

Observation points are positioned in both access (campus) and core (NREN) networks. As opposed to darknets, these networks have seen real IPv6 traffic for several years. Only by actual measurements, we can determine the need for new security approaches.

Access networks

Core networks



**As people have a hard time deploying IPv6 itself, its security tends to come second. It's time to see what is really going on.**

**Many threats, on many levels:**

- Transitional technologies introduce their own attack vectors
- ICMP6 forms a basis of many attacks, but can not simply be blocked
- End users are unaware of connectivity over IPv6

**Flexible, easy to deploy observation points are key:**

- Docker-based, thus transparent and deployable
- Bro-based, allowing both packet and flow-based approaches
- Only local, on-site storage of traffic for further analysis
- No sensitive data leaves the network
- Constantly improving, real-time threat detection



UNIVERSITY  
OF TWENTE.



## C Characterizing and Mitigating the DDoS as a Service Phenomenon

This appendix include the poster titled “Characterizing and Mitigating the DDoS as a Service Phenomenon”, by J. Santanna, A. Sperotto and A. Pras (UT). The poster has been presented at the Traffic Monitoring and Analysis PhD School 2015, and has been awarded a prize in the Centre for Telematics and Information Technology<sup>30</sup> Symposium 2015 [60].

The poster presents a summary of achievements on the DDoS as a service investigation. The poster is divided in five parts. The first depicts the problem statement followed by the explanation on “how Booters work?”. After that two methodologies towards the mitigation of the DDoS as a service phenomenon are shown. Each part of the poster has a QR code pointing to a relevant paper from the same authors. The last part of our poster highlight directions to future work.

---


<sup>30</sup><http://www.utwente.nl/ctit/>

# CTIT

## CHARACTERIZING AND MITIGATING THE DDOS AS A SERVICE PHENOMENON

José Jair Santanna, Anna Sperotto, and Aiko Pras  
{j.j.santanna,a.sperotto,a.pras}@utwente.nl


### The Problem



[by using Booters] even people without enough technical knowledge are able to launch DDOS attacks against anyone in the Internet.



### How Booters work?



Prices starting from USD1; 15 different types of attacks; Clients that launched more than 200 attacks per day;




### "The Bloodhound Methodology"



We found around 1000 Booters and our crawler had 85% of accuracy with potential to work close to 100%.




### "The Baseball Methodology"



We publicly share the collected attack data. We reveal that the volume of attacks measured can increase at least 3000 times.



### Our Batter(s) is Evolving



Hypothesis:

Any Booter can be mitigated by buying, measuring, fingerprinting, and afterwards filtering the attacks.

UNIVERSITEIT TWENTE.

Disclaimer: We are aware that research of this nature may touch on, or cross, legal boundaries, but we are convinced that the results from this research will benefit future mitigation methods and thus help combat Booters, both operationally as well as legally. In order to be transparent about our work, we have informed the office of the public prosecutor in the Netherlands about our intention to pursue this research. We also conduct our research with support of the ethical advisor of the University of Twente.



## **D DDoS Attack Mitigation using OpenFlow-based SDN**

This appendix include the poster titled “DDoS Attack Mitigation using OpenFlow-based SDN”, by M.Jonker and A. Sperotto (UT). The poster has been presented at the Traffic Monitoring and Analysis PhD School 2015 [61]

This poster demonstrates the worrying landscape of DDoS attacks. This has seen not only an increase in occurrence and magnitude of attacks over the last few years, but is also aggravated by the concept of Booters. As the poster further shows, existing mitigation solutions have several issues. They are typically coarse when it comes to Internet traffic diversion, and the diversion in itself is cause for privacy concerns. Software Defined Networking (SDN) is suspected to help in address some of these problems, but SDN comes with some challenges on its own. The poster proposes addressing these challenges, and it sketches an approach to DDoS attack mitigation architecture based on SDN.

# DDoS Attack Mitigation using OpenFlow-based SDN

Mattijs Jonker and Anna Sperotto


Design and Analysis of Communication Systems (DACS)

Centre for Telematics and Information Technology (CTIT)

m.jonker@utwente.nl

## Distributed Denial-of-Service (DDoS) Attacks

... have become an increasing threat, recently reaching traffic volumes of up to 500 Gbps. To make matters worse, Booters allow for the layman to launch attacks in the order of tens of Gbps in exchange for a few euros.




Attacks increase in occurrence & magnitude

Some attack types use core parts of the Internet (e.g., DNS), which makes them nearly impossible to mitigate with strictly on-premise solutions. As a result, a market for mitigation solutions was created, which gave rise to DDoS Protection Service (DPS) providers.

## Software Defined Networking (SDN)


In SDN, the control plane and data plane of the network are decoupled. This has many advantages, such as centralized control over forwarding decisions, dynamic updating of forwarding rules, and easier and more flexible network configuration.





There are challenges when it comes to SDN, such as the performance of programmable forwarding devices, and its susceptibility to attacks.

## DDoS Attack Mitigation


Even though SDN appears well-suited for mitigation, existing solutions (e.g., CloudFlare) rely on techniques such as DNS anycast and BGP. These techniques typically re-route all customer traffic through (remote) scrubbing centres at L3.




These techniques have several disadvantages ...




We propose researching if OpenFlow can be used in a DDoS attack mitigation architecture, and how such an architecture will operate. The research is mostly measurement-based, involving benchmarks of OF-capable devices, comparative measurements with other techniques, etc.



Measurements in terms of performance, flexibility, scalability, advantages of more granular (>L3) re-routing, etc.



Firewall 'as-a-Service', at >L3. Less need for re-routing, decreases load on infrastructure, and reduces unwanted monitoring concerns.



Shift routing control to customer, to steer clear from untrusted third parties, and allow subscription to trusted third parties.



The goal of this research is to evaluate OpenFlow for use in DDoS attack mitigation, and to design and to develop a mitigation architecture.

SURF NET



UNIVERSITY OF TWENTE

## References

- [1] R. van Rijswijk-Deij, A. Sperotto, and A. Pras. DNSSEC and Its Potential for DDoS Attacks: A Comprehensive Measurement Study. In *Proceedings of the 2014 Internet Measurement Conference (IMC 2014)*, pages 449–460, Nov 2014.
- [2] W. de Vries, J.J. Santanna, A. Sperotto, and A. Pras. How Asymmetric Is the Internet? A Study to Support the use of Traceroute. In *Proc. of the 9th International Conference on Autonomous Infrastructure, Management and Security (AIMS 2015)*, pages 113–125. June 2015.
- [3] D. Dönni, G. S. Machado, C. Tsiaras, and B. Stiller. Schengen Routing: A Compliance Analysis. In *9th International Conference on Autonomous Infrastructure, Management, and Security (AIMS 2015), Lecture Notes in Computer Science, Springer*, 2015.
- [4] Roland van Rijswijk-Deij, Mattijs Jonker, Anna Sperotto, and Aiko Pras. The Internet of Names: A DNS Big Dataset. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication (SIGCOMM 2015)*, pages 91–92, 2015.
- [5] O. van der Toorn, R. Hofstede, M. Jonker, and A. Sperotto. A first look at HTTP(S) intrusion detection using NetFlow/IPFIX. In *Integrated Network Management (IM), 2015 IFIP/IEEE International Symposium on*, pages 862–865, May 2015.
- [6] M. Jonker, R. Hofstede, A. Sperotto, and A. Pras. Unveiling flat traffic on the Internet: An SSH attack case study. In *Integrated Network Management (IM), 2015 IFIP/IEEE International Symposium on*, pages 270–278, May 2015.
- [7] J. van der Hooft, S. Petrangeli, M. Claeys, J. Famaey, and F. De Turck. A Learning-Based Algorithm for Improved Bandwidth-Awareness of Adaptive Streaming Clients. In *International Symposium on Integrated Network Management (IM 2015)*, pages 131–138. IEEE, 2015.
- [8] J.J. Chromik, J.J. Santanna, A. Sperotto, and A. Pras. Booter websites characterization: Towards a list of threats. In *Proc. of the XXXIII Simpósio Brasileiro de Redes de Computadores e Sistemas Distribuídos (SBRC 2015)*, May 2015.
- [9] G. Hurel, R. Badonnel, A. Lahmadi, and O. Festor. Towards Cloud-Based Compositions of Security Functions For Mobile Devices. In *IFIP/IEEE International Symposium on Integrated Network Management (IM'15)*, 2015.
- [10] K. Scarfone and P. Mell. Guide to Intrusion Detection and Prevention Systems (IDPS). *NIST Special Publication*, 800(2007):94, 2007.
- [11] M. Golling and B. Stelte. Requirements for a Future EWS-Cyber Defence in the Internet of the Future. In *Cyber Conflict (ICCC), 2011 3rd International Conference on*, 2011.
- [12] GÉANT. Breakthrough GÉANT Network Marks Ten Years of Success: High Bandwidth pan-European Research Network Continues Advances with 100 Gbps Plans , November 2010.
- [13] Mario Golling, Rick Hofstede, and Robert Koch. Towards Multi-layered Intrusion Detection in High-Speed Backbone Networks. In *Proceedings of the NATO CCD COE 6th International Conference on Cyber Conflict, CyCon'14*, 2014.
- [14] Rick Hofstede, Václav Bartoš, Anna Sperotto, and Aiko Pras. Towards Real-Time Intrusion Detection for NetFlow/IPFIX. In *Proceedings of the 9th International Conference on Network and Service Management, CNSM'13*, pages 227–234, 2013.



- [15] L. Hellemons, L. Hendriks, R. Hofstede, A. Sperotto, R. Sadre, and A. Pras. SSHCure: A Flow-Based SSH Intrusion Detection System. In *Proc. of the 6th International Conference on Autonomous Infrastructure, Management and Security (AIMS 2012)*, volume 7279, pages 86–97. IEEE, 2012.
- [16] R. Koch, M. Golling, and G. Dreo Rodosek. Evaluation of State of the Art IDS Message Exchange Protocols. In *International Conference on Communication and Network Security (CNS 2013)*, 2013.
- [17] O. Festor, A. Lahmadi, R. Hofstede, and A. Pras. Information Elements for IPFIX Metering Process Location (Internet Draft). <http://tools.ietf.org/html/draft-irtf-nmrg-location-ipfix-04>, July 2015.
- [18] R. J. Hofstede and T. Fioreze. SURFmap: A Network Monitoring Tool Based on the Google Maps API. In *IFIP/IEEE International Symposium on Integrated Network Management (IM 2009)*, June 2009.
- [19] Christos Tsiaras, Anuj Sehgal, Sebastian Seeber, Daniel Dönni, Burkhard Stiller, Jürgen Schönwälder, and Gabi Dreo Rodosek. Towards evaluating type of service related quality-of-experience on mobile networks. In *7th IFIP Wireless and Mobile Networking Conference (WMNC), Vilamoura, Algarve, Portugal, May 2014.*, 2014.
- [20] Google Play: Bonafide+ Application. <https://play.google.com/store/apps/details?id=de.jacobs.university.cnds.bonafide.plus>. Accessed: 2014-09-22.
- [21] Measurement Lab. <http://www.measurementlab.net/>. Accessed: 2014-09-22.
- [22] Emanics Lab. <http://www.emanicslab.org>. Accessed: 2014-09-22.
- [23] Bondafide. <http://www.bonafide.pw>. Accessed: 2014-09-22.
- [24] Susan Landau. Making sense from snowden: What's significant in the nsa surveillance revelations. *IEEE Security & Privacy*, (4):54–63, 2013.
- [25] John Blau. Nsa surveillance sparks talk of national internets. *Spectrum, IEEE*, 51(2):14–16, 2014.
- [26] European Union Commission. The Schengen Area. [http://biblio.ucv.ro/bib\\_web/bib\\_pdf/EU\\_books/0056.pdf](http://biblio.ucv.ro/bib_web/bib_pdf/EU_books/0056.pdf). Accessed in July, 2015.
- [27] Norbert Pohlmann, Michael Sparenberg, Illya Siromaschenko, and Kilian Kilden. Secure communication and digital sovereignty in europe. In *ISSE 2014 Securing Electronic Business Processes*, pages 155–169. Springer, 2014.
- [28] RIPE Network Coordination Center. RIPE ATLAS. <http://atlas.ripe.net>. Accessed in July, 2015.
- [29] Maxmind. GeoLite Legacy Downloadable Databases. <http://dev.maxmind.com/geoip/legacy/geolite/>. Accessed in July, 2015.
- [30] Computerwoche.: Internet-Verband ECO Beklagt Scheindiskussion um Schengen-Routing. <http://www.computerwoche.de/a/internet-verband-eco-beklagt-scheindiskussion-um-schengen-routing,2556658>. Accessed in July, 2015.



- [31] Yihua He, Michalis Faloutsos, Srikanth Krishnamurthy, and Bradley Huffaker. On routing asymmetry in the internet. In *Global Telecommunications Conference, 2005. GLOBECOM'05. IEEE*, volume 2, pages 6–pp. IEEE, 2005.
- [32] V. Bajpai and J. Schönwälder. A Survey on Internet Performance Measurement Platforms and Related Standardization Efforts. *Communications Surveys & Tutorials, IEEE*, 17(3):1313–1341, 2015.
- [33] Vaibhav Bajpai, Steffie Jacob Eravuchira, and Jürgen Schönwälder. Lessons Learned From Using the RIPE Atlas Platform for Measurement Research. *SIGCOMM Comput. Commun. Rev.*, 45(3):35–42, July 2015.
- [34] Aemen Lodhi, Natalie Larson, Amogh Dhamdhere, Constantine Dovrolis, and kc claffy. Using peeringDB to Understand the Peering Ecosystem. *SIGCOMM Comput. Commun. Rev.*, 44(2):20–27, April 2014.
- [35] Florian Weimer. Passive DNS Replication. In *Proc. of the 17th FIRST Conference (FIRST 2005)*, 2005.
- [36] J. J. Cardoso de Santanna, R. M. van Rijswijk-Deij, A. Sperotto, M. Wierbosch, L. Zambenedetti Granville, and A. Pras. Booters - an analysis of DDoS-as-a-Service attacks. In *Proceedings of the IFIP/IEEE International Symposium on Integrated Network Management (IM 2015)*, pages 243 –251, May 2015.
- [37] Prolexic. Prolexic Quarterly Global DDoS Attack Report (Q1 2014). <http://www.prolexic.com/knowledge-center-ddos-attack-report-2014-q1.html>.
- [38] J. Damas, M. Graff, and P. Vixie. Extension Mechanisms for DNS (EDNS(0)). RFC 6891, 2013.
- [39] J. Postel. Character Generator Protocol. RFC 689, 1983.
- [40] Roland van Rijswijk-Deij, Anna Sperotto, and Aiko Pras. Making the Case for Elliptic Curves in DNSSEC. *SIGCOMM Comput. Commun. Rev.*
- [41] G. van den Broek, R. van Rijswijk, A. Sperotto, and A. Pras. DNSSEC Meets Real World: Dealing with Unreachability Caused by Fragmentation. *IEEE Communications Magazine*, 52(April):154–160, 2014.
- [42] R. van Rijswijk-Deij, A. Sperotto, and A. Pras. DNSSEC and its potential for DDoS attacks. In *ACM IMC 2014*, Vancouver, BC, Canada, 2014.
- [43] T. Winter, P. Thubert, A. Brandt, J. Hui, R. Kelsey, P. Levis, K. Pister, R. Struik, J. Vasseur, and R. Alexander. RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks. *IETF RFC 6550*, March 2012.
- [44] T. Phinney, P. Thubert, and R. A. Assimiti. RPL Applicability in Industrial Networks. *IETF I-D <draft-ietf-roll-rpl-industrial-applicability-02>*, October 2013.
- [45] A. Brandt, E. Baccelli, R. Cragie, and P. van der Stok. Applicability Statement: The use of the RPL protocol suite in Home Automation and Building Control. *IETF I-D <draft-ietf-roll-applicability-home-building-06>*, December 2014.
- [46] D. Popa, M. Gillmore, L. Toutain, J. Hui, R. Ruben, and K. Monden. Applicability Statement for the Routing Protocol for Low Power and Lossy Networks (RPL) in AMI Networks. *IETF I-D <draft-ietf-roll-applicability-ami-09>*, July 2014.

- [47] J. Hui and J. Vasseur. The Routing Protocol for Low-Power and Lossy Networks (RPL) Option for Carrying RPL Information in Data-Plane Datagrams. IETF RFC 6553, March 2012.
- [48] Philip Alexander Levis, Neil Patel, David Culler, and Scott Shenker. Trickle: A Self Regulating Algorithm for Code Propagation and Maintenance in Wireless Sensor Networks. In *1st Symposium on Networked Systems Design and Implementation (NSDI)*, San Francisco, CA, USA, March 2004.
- [49] Anth  a Mayzaud, Anuj Sehgal, R  mi Badonnel, Isabelle Chrisment, and J  rgen Sch  nw  lder. Mitigation of topological inconsistency attacks in rpl-based low-power lossy networks. *International Journal of Network Management*, 2015.
- [50] Anuj Sehgal, Anth  a Mayzaud, R  mi Badonnel, Isabelle Chrisment, and J  rgen Sch  nw  lder. Addressing DODAG Inconsistency Attacks in RPL Networks. In *Proc. of GIIS conference*, 2014.
- [51] Cisco Systems Inc. How To Calculate Bandwidth Utilization Using SNMP. [http://www.cisco.com/image/gif/paws/8141/calculate\\_bandwidth\\_snmp.pdf](http://www.cisco.com/image/gif/paws/8141/calculate_bandwidth_snmp.pdf), 2005. Online. Accessed May 2014.
- [52] Milosz Marian Hulboj and Ryszard Erazm Jurga. CERN Investigation of Network Behaviour and Anomaly Detection. In *Proceedings of the 12th International Symposium on Recent Advances in Intrusion Detection*, RAID, pages 353–354, 2009.
- [53] Elisa Jasinska. sFlow: I can feel your traffic. In *Proceedings of the 23rd Chaos Communication Congress*, 23C3, pages 1–8, 2006.
- [54] Peter Phaal, Sonia Panchen, and Neil McKee. InMon Corporation’s sFlow: A Method for Monitoring Traffic in Switched and Routed Networks. RFC 3176, 2001.
- [55] Tanja Zseby, Maurizio Molina, Nick Duffield, Saverio Niccolini, and Frederic Raspall. Sampling and Filtering Techniques for IP Packet Selection. RFC 5475, 2009.
- [56] Hans van den Berg, Michel Mandjes, Remco van de Meent, Aiko Pras, Frank Roijers, and Pieter Venemans. QoS-aware bandwidth provisioning for IP network links. *Elsevier Computer Networks*, 50(5):631–647, 2006.
- [57] Aiko Pras, Lambert J. M. Nieuwenhuis, Remco van de Meent, and Michel R. H. Mandjes. Dimensioning Network Links: A New Look at Equivalent Bandwidth. *IEEE Network*, 23(2):5–10, 2009.
- [58] Ricardo de O. Schmidt, Ramin Sadre, Anna Sperotto, Hans van den Berg, and Aiko Pras. Impact of Packet Sampling on Link Dimensioning. *IEEE TNSM*, 12(3):392–405, 2015.
- [59] L. Hendriks, A. Sperotto, and A. Pras. Determining the State of Security in the IPv6 Internet. TERENA Networking Conference 2015 (TNC 2015), June 2015.
- [60] J.J. Cardoso de Santanna, A. Sperotto, and A. Pras. Characterizing and Mitigating the DDoS as a Service Phenomenon. Traffic Monitoring and Analysis PhD School (TMA 2015), April 2015.
- [61] M. Jonker and A. Sperotto. DDoS Attack Mitigation using OpenFlow-based SDN. Traffic Monitoring and Analysis PhD School (TMA 2015), April 2015.