

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

WP5 — Network and Service Monitoring

Deliverable D5.2 — Second 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: D5.2 — Second year deliverable on network and service monitoring
Type: Public
Editor(s): Anna Sperotto
E-mail: a.sperotto@utwente.nl
Doc ID: D5.2
Delivery Date: 31.10.2014
Author(s): Anthea Mayzaud, Anuj Sehgal, Abdelkader Lahmadi,
Christos Tsiaras, Daniel Dönni, Daphne Tuncer
Marinos Charalambides, Mario Flores, Jeroen Famaey,
Mario Golling, Maxim Claeys, Niels Bouten, Gaëtan Hurel
Radhika Garg, Rashid Mijumbi, Ricardo Schmidt, Frank Tietze,
Rick Hofstede, Sebastian Seeber, Corinna Schmitt,
Guilherme Sperb Machado, Jair Santanna, Stefano Petrangeli

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 D5.2 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 Y2. Then, we report a set of selected research highlights, in 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 projects and the scientific output for Y2, 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 one, which has been already targeted in Y1. Particularly active and successful is the research conducted in the context of Objective 5 (“To collect (anonymized) monitoring data”) and Objective 6 (“To create annotated traces to assess the quality of different Intrusion Detection Systems”). The research conducted for fulfilling these objectives resulted in some of the top publications for WP5.

The key achievements for Y2 are:

- Following one of the comments received during the Y1 review meeting, in Y2 WP5 focused on publishing in top conferences and journals in the fields of networking, network measurements and network and service management (Sections 5–9). This resulted in papers at, for example, INFOCOM 2014, IMC 2014 and CCR. The paper published at INFOCOM analyses 5 years of network traces from the Japanese National Research and Education Network WIDE (Section 5). The IMC 2014 paper investigates the potential of DNSSEC for abuse in Distributed Denial of Service attacks based on a dataset covering roughly 70% of all registered domains worldwide (Section 8). The CCR paper proposes the first flow-based detection method able to identify successful SSH attacks, and it validates the approach on a campus network, a large European backbone and the a public SSH blacklist (Section 6).
- In Y1, the project published 37 papers. In Y2, we exceeded the previous year number of publications but publishing 50 papers, many of which containing a strong measurement component. Of the published papers, for WP5 and WP6, which also in Y2 are working in close collaboration, 9 are the result of Ph.D. collaborations, and 20 have been achieved in collaboration with other EU projects and institutions. An overview of the scientific output for WP5 is given in Section 2.1. The complete list of FLAMINGO publications can be found in D8.2.
- Collaborations are a core characteristic of the work in WP5. First of all, the Ph.D. collaborations are playing also in Y2 a crucial role in the collaboration between the research work packages. In Y2 we have seen several new collaboration starting, and other coming to fruition with a number of publications, as mentioned above. The joint PhD collaborations are reported in both D5.2 and D6.2 by highlighting the contributions to each work package and its objectives (Section 3). In addition, prototypes and open-source tools are a crucial component for some of the research carried out in this WP. For this reason, WP5 is also closely collaborating with WP1. Finally, WP5 is also collaborating with other European projects and institutions. The activities and output of such collaborations is reported in D3.3 and Section 2.1, respectively.

Contents

1	Introduction	1
2	Objectives and Tasks	2
2.1	S.M.A.R.T. Objectives	2
2.2	Workpackage Objectives	6
2.2.1	Ongoing Objectives	6
2.2.2	Open Objectives	9
2.3	Tasks and Objectives Mapping	9
3	Integration of PhD Students	10
3.1	PhD Student Collaborations	10
3.2	Description of the Collaborations	12
3.2.1	Linking Network Usage Patterns to Traffic Gaussianity Fit (JUB-UT-Pattern) .	13
3.2.2	Energy-aware Traffic Management (UCL-UT-Man)	14
3.2.3	Intrusion Detection Systems (UT-UniBwM-IDS)	15
3.2.4	Towards A Trust Computing Architecture for RPL in Cyber Physical Systems (UniBwM-JUB-RPL)	18
3.2.5	QoE-Driven In-Network Optimization for Adaptive Video Streaming Based on Packet Sampling Measurements (iMinds-UT-QoS)	18
3.2.6	Security of RPL Networks (INRIA-JUB-RPL)	19
3.2.7	Distributed Monitoring Architecture for the Internet of Things (INRIA-JUB-Distr)	20
3.2.8	Mobile Measurements (UZH-JUB-UniBwM-M2)	20
3.2.9	Cache Management (UCL-iMinds-Cache)	20
3.2.10	Management of Virtualized Networks (iMinds-UPC-NetVirt)	21
3.2.11	Cloud Security (INRIA-UniBwM-Cloud)	23
3.2.12	Flowoid: a NetFlow/IPFIX Probe for Android-based Devices (UT-INRIA-Flowoid)	23
4	Data Collection	25
4.1	Data Collection and Joint Security Lab	25
4.2	Publicly-released Datasets	25
5	A Longitudinal Analysis of Internet Rate Limitations	27
5.1	Dataset and Preprocessing	27
5.2	Flow Classification	28
5.3	Highlights of Analysis Results	29

6 SSH Compromise Detection using NetFlow/IPFIX	31
6.1 The detection model	31
6.2 The detection algorithm	32
6.3 Validation results	33
7 DNSSEC Meets Real World: Dealing with Unreachability Caused by Fragmentation	34
7.1 DNSSEC and Fragmentation	34
7.2 Solutions to Fragmentation	35
7.2.1 Avoiding Fragmentation in General	35
7.2.2 Selectively Avoiding Response Fragmentation	35
7.3 Validation Results	35
8 DNSSEC and Its Potential for DDoS attacks	37
8.1 DNS Amplification	37
8.2 Measurement Methodology	37
8.3 Data Sets	38
8.3.1 Source Data	39
8.3.2 Collected data	39
8.4 Highlights of Results	40
9 Link Dimensioning	41
10 Integration of EU research	44
10.1 International Activities	44
10.2 Collaborations with Other EU Projects and Institutions	45
11 Conclusions	46
A Internet Traffic Statistics	49
B Towards Comparability of Intrusion Detection Systems: New Data Sets	51
C SSHCure: SSH Intrusion Detection using NetFlow and IPFIX	53
D SSHCure: SSH Intrusion Detection using NetFlow and IPFIX	55

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 Y2, WP5 has continued and improved research in the field of network and service monitoring, 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 Y2, some coming to maturity after the work carried on in Y1, others starting in Y2. 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 the first year, the research work packages have published 50 papers, both at major conferences and in journals. Several of those papers (10) are the outcome of collaborations within the consortium. Many of the published papers have a measurement component that has been investigated in the context of WP5. In answer to the first review meeting, WP5 has also targeted high-ranking publication venues in the areas of networking, network monitoring as well as network management, which resulted in publications in conferences as IEEE INFOCOM 2014 and IMC 2014, and in journals such as ACM Computer Communication Review, Elsevier Computer Networks (see also Sections 5–9) and IEEE Surveys & Tutorials. In addition, several other papers are currently under review. 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 D8.1.

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, also Y2 has indicated that large scale measurements and measurements of security events are two of the core areas of interested for WP5. The highlights of research focus therefore on these topics. In particular, Section 4 reports WP5's efforts in the area of data collection and the release of publicly available datasets; Section 5 reports on the an in-depth, measurement-based analysis of the Internet rate limitations; Section 6 describes how we can achieve flow-based compromise detection for the SSH protocol, a research effort that has been integrated in SSHCure, an open-source tool for intrusion detection and fingerprinting that has been developed in collaboration with WP1; Sections 7 and 8 analyze DNSSEC and its implication in network security; finally, Section 9 presents a modeling-based approach to link dimensioning that has been validated using traffic traces worldwide.

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 Y2. 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 year of the project, seven Ph.D. students have joined FLAMINGO as *fully integrated Ph.D. students*. In the second year, 7 more Ph.D. students have joined the NoE. These students, their affiliation and the co-supervising institution are listed in D8.2.

For the FLAMINGO project, collaboration is at the basis of research. Therefore, Ph.D. students are not working in isolation, but they are encouraged to collaborate with other institutions. For this reason, there is not a one-to-one match between a Ph.D. student and a single WP. In addition, it is important to highlight that Ph.D. collaborations are taking place not only among fully integrated Ph.D. students, but also with 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. In addition,

- **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 Y1, the project fulfilled and exceeded the expected number of publications. In Y2, the research work packages have published 50 papers, both at major conferences and in journals. In addition, several other papers are currently under review. The complete list of published papers can be found in D8.2. Given the strong inter-dependence between monitoring (WP5) and configuration and repair actions (WP6), as described in the DoW (Sec. B.1.1.2), we highlights in this deliverable, however, the scientific output of WP5 and WP6 with respect to the collaboration with other EU projects and institutions and within the FLAMINGO consortium Tables 1–3. Last, according to the suggestions of the reviewers in Y1 evaluation, 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 published at the IEEE INFOCOM 2014, a paper published at Internet Measurement Conference (IMC 2014), a publication in the ACM SIGCOMM Computer Communication Review (CCR), one in IEEE Communications Magazine, one at Elsevier Computer Networks and one in IEEE Surveys & Tutorials.

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

Authors	Title	Venue	EU project/ institution
R. de O. Schmidt, R. Sadre, A. Sperotto, H. van den Berg, and A. Pras	A hybrid procedure for efficient link dimensioning	Elsevier Computer Networks	Aalborg University
R. de O. Schmidt, R. Sadre, N. Melnikov, J. Schönwälder, and A. Pras.	Linking network usage patterns to traffic Gaussianity fit	IFIP Networking 2014	Aalborg University
F. Francois, N. Wang, K. Moessner, S. Georgoulas, and R. de O. Schmidt	Leveraging MPLS Backup Paths for Distributed Energy-Aware Traffic Engineering.	TNSM	University of Bristol, University of Surrey
I. Drago, R. de O. Schmidt, R. J. Hofstede, A. Sperotto, M. Karimzadeh, B. R. H. M. Haverkort, and A. Pras	Networking for the Cloud: Challenges and Trends	PIK - Praxis der Informationsver- arbeitung und Kommunikation	Politecnico di Torino, Mobile Cloud Networking
R. Hofstede, P. Celeda, B. Trammell, I. Drago, R. Sadre, A. Sperotto, and A. Pras	Flow Monitoring Explained: From Packet Capture to Data Analysis with NetFlow and IPFIX	IEEE Communications Surveys & Tutorials	Masaryk University, ETH, Politecnico di Torino, Aalborg University
G. van den Broek, R. van Rijswijk-Deij, A. Sperotto, and A. Pras	DNSSEC meets real world: dealing with unreachability caused by fragmentation	IEEE Communications Magazine	SURFnet BV
R. van Rijswijk-Deij, A. Sperotto, and A. Pras	DNSSEC and Its Potential for DDoS Attacks	ACM IMC 2014	SURFnet BV
C. Tsiaras, A. Sehgal, S. Seeber, D. Dönni, B. Stiller, J. Schönwälder, and G. Dreo Rodosek	Towards Evaluating Type of Service Related Quality-of-Experience on Mobile Net- works	WMNC 2014	SmartenIT
C. Tsiaras, S. Liniger, and B. Stiller	Automatic and on-demand Mobile Network Operator (MNO) selection mechanism demonstration	NOMS 2014 (Demo paper)	SmartenIT
C. Tsiaras, S. Liniger, and B. Stiller	An automatic and on-demand MNO selection mechanism	NOMS 2014	SmartenIT
C. Tsiaras, and B. Stiller	A Deterministic QoE Formalization of User Satisfaction Demands (DQX)	LCN 2014	SmartenIT
R. Mijumbi, J.L. Gorricho, and J. Serrat	Contributions to Efficient Resource Management in Virtual Networks	AIMS 2014	EVANS

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

Authors	Title	Venue	EU project/ institution
R. Mijumbi, J. Serrat, J.L. Gorricho, M. Claeys, F. De Turck, and S. Latré	Design and evaluation of learning algorithms for dynamic resource management in virtual networks	NOMS 2014	EVANS, University of Antwerp
R. Mijumbi, J.L. Gorricho, J. Serrat, M. Claeys, J. Famaey, and F. De Turck	Neural Network-based Autonomous Allocation of Resources in Virtual Networks	EUCNC 2014	EVANS
M. Claeys, S. Latré, J. Famaey, and F. De Turck	Design and evaluation of a self-learning http adaptive video streaming client	IEEE Communications Letters	University of Antwerp
N. Bouten, M. Claeys, S. Latré, J. Famaey, W. Van Leekwijck, and F. De Turck	Deadline-based Approach for Improving Delivery of SVC-based HTTP Adaptive Streaming Content	QCMan 2014	University of Antwerp
J. T. Araújo, R. Landa, R. G. Clegg, and G. Pavlou	Software-defined network support for transport resilience	NOMS 2014	ALIEN
J. T. Araújo, R. Landa, R. Clegg, G. Pavlou, and K. Fukuda	A longitudinal analysis of Internet rate limitations	INFOCOM 2014	NII International Internship Program
P. Porambage, C. Schmitt, A. Gurtov, and S. Gerdes	PAuthKey: A Pervasive Authentication Protocol and Key Establishment Scheme for Wireless Sensor Networks in Distributed IoT Applications	International Journal of Distributed Sensor Networks	SmartenIT, Oulu University and Aalto University
M. Claeys, D. Tuncer, J. Famaey, M. Charalambides, S. Latré, F. De Turck, and G. Pavlou	Towards Multi-tenant Cache Management for ISP Networks	EUCNC 2014	University of Antwerp
M. Karimzadeh, A. Sperotto, and A. Pras	Software Defined Networking to Improve Mobility Management Performance	AIMS 2014	Mobile Cloud Networking
C. Schmitt, T. Kothmayr, B. Ertl, W. Hu, L. Braun, and G. Carle	Tiny IPFIX: An efficient application protocol for data exchange in cyber physical systems.	Elsevier Computer Communications	TU Munich and TU Berlin

Table 3: Publications authored by multiple FLAMINGO partners.

Authors	Title	Venue	FLAMINGO partners
R. de O. Schmidt, R. Sadre, N. Melnikov, J. Schönwälder, and A. Pras	Linking network usage patterns to traffic Gaussianity fit	IFIP Networking 2014	UT, JUB
A. Mayzaud, A. Sehgal, R. Badonnel, I. Chrisment, and J. Schönwälder	A Study of RPL DODAG Version Attacks	AIMS 2014 Best Paper Award	INRIA, JUB
C. Tsiaras, A. Sehgal, S. Seeber, D. Dönni, B. Stiller, J. Schönwälder, and G. Dreo Rodosek	Towards Evaluating Type of Service Related Quality-of-Experience on Mobile Networks	WMNC 2014	UZH, JUB, UniBwM
A. Sehgal, A. Mayzaud, R. Badonnel, I. Chrisment, and J. Schönwälder	Addressing DODAG Inconsistency Attacks in RPL Networks	GIIS 2014	JUB, INRIA
M. Golling, R. Hofstede, and R. Koch	Towards Multi-layered Intrusion Detection in High-Speed Backbone Networks	CyCon 2014	UniBwM, UT
R. Mijumbi, J. Serrat, J.L. Gorricho, M. Claeys, F. De Turck, and S. Latré	Design and evaluation of learning algorithms for dynamic resource management in virtual networks	NOMS 2014	UPC, iMinds
R. Mijumbi, J.L. Gorricho, J. Serrat, M. Claeys, J. Famaey, and F. De Turck	Neural Network-based Autonomous Allocation of Resources in Virtual Networks	EUCNC 2014	UPC, iMinds
M. Claeys, D. Tuncer, J. Famaey, M. Charalambides, S. Latré, F. De Turck, and G. Pavlou	Towards Multi-tenant Cache Management for ISP Networks	EUCNC 2014	iMinds, UCL
M. Golling, R. Koch, P. Hillmann, R. Hofstede, and F. Tietze	YANG2UML: Bijective Transformation and Simplification of YANG to UML	CNSM 2014	UniBwM, UT

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 on in this WP (*open objectives*).

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. A notable one is the Dagstuhl Seminar *Ethics in Data Sharing*¹, where WP5 members have been active, both in the organization (A. Pras) as well as in participating. The seminar has taken place from January 26 to January 31, 2014. Additional information can be found in D3.3. The Dagstuhl seminar has set the basis for the “Legal and Ethical Facets of Data Sharing” scenario in WP7. Another important activity has been the organization of the 8th International Conference on Autonomous Infrastructure, Management and Security (AIMS 2014), also described in D3.3. Section 10 provides more information regarding the activities for this objective.

Objective 2: To create and maintain articles within Wikipedia and other online systems in this area – The research conducted in Y1 has allowed us to generate valuable knowledge that can be used for contributing to Wikipedia. In collaboration with WP2, WP5 has identified a set of Wikipedia articles where a contribution would be beneficial. WP5 has for example contributed to pages as *Denial of Service attack*², *NetFlow*³, *IP Flow Information Export*⁴, *sFlow*⁵, *Software Defined Networking*⁶ and *Internet of Things*⁷. For more information on this topic, we refer the reader to D2.2.

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 the architecture development has been carried on by WP6, while WP5 has supported the architecture when necessary with measurement. For more details about the architecture, we refer the reader to D6.2.

Objective 5: To collect (anonymized) monitoring data – Several data collection activities are taking place among FLAMINGO partners. The *Internet Traffic Statistics* project, introduced in Y1, is currently collecting weekly traffic reports from supporting operators. Also, new contacts with other operators have been made during the TERENA Networking Conference 2014. Appendix A provides additional details on the project.

¹<http://www.dagstuhl.de/en/program/calendar/semhp/?semnr=14052>

²http://en.wikipedia.org/wiki/Denial-of-service_attack

³<http://en.wikipedia.org/wiki/NetFlow>

⁴http://en.wikipedia.org/wiki/IP_Flow_Information_Export

⁵<http://en.wikipedia.org/wiki/SFlow>

⁶http://en.wikipedia.org/wiki/Software-defined_networking

⁷http://en.wikipedia.org/wiki/Internet_of_Things

Focusing on security, annotated flow data traces have been created with the specific goal of evaluating Intrusion Detection Systems. Objective 6 will provide additional information about these activities. Attack traces for analysis of Booters and DDoS attacks have been collected at the UT campus, and a new data collection is currently being deployed at the Amsterdam Internet Exchange in collaboration with SURFnet.

A research topic that has been intensively data-driven focus on DNS(SEC), a research conducted in collaboration between UT and SURFnet. A first research topic concerns the effects of packet fragmentation on DNSSEC resolvers. A paper on the topic has been published in IEEE Communication Magazine [1] (see also Section 7). A second research topic focuses instead on the implications of DNSSEC in security, as it quantifies the effect of DNS amplification in DNSSEC and its possible effect on DDoS attacks. A paper on this research has been published at the Internet Measurement Conference 2014 (IMC 2014). The dataset associated with this research has been made publicly available in anonymized form and it is available at <http://www.simpleweb.org/wiki/Traces>. More information about this research is available in Section 8.

Another research topic where monitoring data plays a fundamental role is the research on link dimensioning carried on by R. Schmidt at UT. Recently, this line of research led to a publication in the Elsevier journal *Computer Networks*. More information about this research is available in Section 9. UT is also investigating the possibility of using SDN protocols, specifically OpenFlow, as a way of obtaining flow-like measurements.

Finally, although for this research the authors have made use of measurements conducted by third parties, we mention here also the study on TCP congestion control mechanisms and a longitudinal analysis of Internet rate limitation, performed by UCL and published at INFOCOM 2014 (see Section 5).

More information about the data collection activities carried on in WP5 can be found in Section 4.

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 Y2. 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.

UT has continued the development of SSHCure, an open-source flow-based IDS developed by R. Hofstede and L. Hendriks. In particular, the researcher have enhanced the *compromise detection* capability of SSHCure (i.e., the tool aims at identifying not just brute force attacks, but specifically *successful* attacks that result in a compromise), in this way advancing the state-of-the-art in the field. The tool includes domain knowledge about the specific behavior of attack tools, with the goal of reducing the number of false positives. The tool has been validated based on flow data traces and server logs from the UT network, flow data traces from the CESNET network and the OpenBL SSH blacklist. The datasets are available in anonymized form at <http://www.simpleweb.org/wiki/Traces>. The research on this topic resulted in a paper published in ACM Computer Communication Review [2]. Additional information about this research can be found in Section 6.

In addition, UT has investigated how to improve flow-based SSH brute force attack detection in the presence of so-called *non-flat traffic*. The rationale behind this approach is that automated attack traffic, such as the one in a brute force attack, will show several flows with (close to) equal characteristics, called flat traffic. However, TCP retransmissions and control information are typically included in the flow metering process and accounted for in the packets and byte counters of exported flows. The research of R. Hofstede and M. Jonker (M.Sc. at the UT) proved that removing TCP retransmission and TCP control traffic from SSH flows

reveals flat traffic patterns and improves the detection of SSHCure of up to 20%. The method was validated using traffic captures from UT and CESNET. A paper on the topic has been submitted to the IM 2015 conference.

Last, UT has also worked in developing a targeted detection scheme for HTTP(S)-based brute force attacks, which target web applications with the intention of gaining unauthorized access to password-protected pages. Such a research has been carried on by R. Hofstede and O. van der Toorn (B. Sc. student at UT) and it focussed on analyzing malicious code attack patterns against commonly used content management systems (e.g., Joomla, Wordpress and Drupal) and extracting appropriate flow-level signatures. The approach has been evaluated based on HTTP flow traces of relevant web servers captured at the UT campus network. A paper on this topic has been submitted to the IM 2015 conference.

An initiative towards the creation of new labeled traces for comparability of intrusion detection systems is currently carried on at UniBwM. A poster on the topic has been presented at the TERENA Networking Conference 2014 (see also Appendix B).

Objective 7: To investigate the applicability of different AI and machine learning techniques for flow analysis

– Several activities on the topic of AI and machine learning have started in Y2. An example is given by the research on HTTP adaptive streaming carried on by iMinds [3]. HTTP Adaptive Streaming (HAS) is becoming the de facto standard for Over-The-Top-based video streaming services such as YouTube and Netflix. By splitting a video into multiple segments of a couple of seconds and encoding each of these at multiple quality levels, HAS allows a video client to dynamically adapt the requested quality during the playout to react to network changes. However, state-of-the-art quality selection heuristics are deterministic and tailored to specific network configurations. Therefore, they are unable to cope with a vast range of highly dynamic network settings. In [3], a novel Reinforcement Learning-based HAS client is presented and evaluated. The self-learning HAS client dynamically adapts its behaviour by interacting with the environment to optimize the Quality of Experience (QoE), the quality as perceived by the end-user. The proposed client has been thoroughly evaluated using a network-based simulator and is shown to outperform traditional HAS clients by up to 13% in a mobile network environment.

UT, in collaboration with the Hochschule Darmstadt (Prof. H. Baier), is currently investigating how machine learning can be applied to network security and, more specifically, to Botnet detection. The Ph.D. student C. Dietz (Hochschule Darmstadt and UT) is currently developing a machine learning toolbox that will in the future allow to identify which approach is more suitable for the specific problem of Botnet detection.

A collaboration on the topic of machine learning is also currently starting between iMinds and UT. The goal of such a collaboration is a to create a structured overview of the state of the art with the goal of identify the most suitable machine learning approached for a set of specific network management application areas. A paper on the topic will be submitted to IEEE Surveys and Tutorials.

Objective 8: To propose novel solutions for intrusion detection and fingerprinting

– Also in Y2, Intrusion Detection and fingerprinting remained one of the core focus of WP5. Several activities that contribute to this objective have been described in Objective 6. We therefore refer the reader to Objective 6 for research activities related to SSH and HTTP(S) flow-based intrusion detection. In addition, in Y2 WP5 has also focused on the topic of on probe real-time flow-based intrusion detection, with the specific goal of developing a detection mechanism that can be deployed on widely-adopted Cisco routers, in particular the Cisco Catalyst 6500, since this is one of the most often deployed. The research, carried on by R. Hofstede and D. van der Steeg (M.Sc. student at UT) has brought operational experience on how this type

of systems can support intrusion detection and has highlighted the possible pitfalls of such approach. A paper on this topic has been submitted to the IM 2015 conference.

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 how clouds systems can be used to offload security mechanisms for mobile devices. The rationale behind this research is that dedicated security applications should be devoted to the task of preventing smartphones and tablets from being compromised. However, these applications may have a significant impact on the device resources. A new approach is the outsourcing of costly security operations to the cloud, where they can be dynamically activated and configured on demand. A paper providing the research plan for this topic has been accepted at AIMS 2014 [4].

A second line of research focuses on the link between cloud data centers and network monitoring. Cloud providers have repeatedly been related to reports of major failures, including outages and performance degradation. The internal network of cloud data centers has frequently been identified as a root-cause of these problems, showing that network provisioning and monitoring is still a major challenge for the deployment of cloud services. A positioning paper on the topic of *networking for the cloud* has been published in the “Praxis der Informationsverarbeitung und Kommunikation” [5]. This paper argues that today’s technologies for measuring and monitoring Internet traffic could be applied in the context of the internal network of cloud data centers as well. Also, the paper shows the suitability of flow-based traffic measurements for monitoring cloud services. Then, it presents a case on bandwidth capacity provisioning to exemplify how flow-based measurements can be used to guarantee the performance of cloud services. Finally, it discusses future directions in the development of new cloud services. The paper advocates that next-generation cloud services will not only rely on the Internet as a means to reach users, but also influence how the Internet itself is organized.

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 worked carried on 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”.

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.

Following the suggestion received after the first year review, 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 4: Mapping of objectives and tasks.

Objective	Task 5.1	Task 5.2	Task 5.3	Status	Additional Material
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	D2.2
WP Objective 3 Monitoring Architecture	X			Ongoing	D6.2
WP Objective 4 Flow Query Language			X	Open	
WP Objective 5 Monitoring Data		X		Ongoing	Section 4 Section 7 Section 8 Section 9
WP Objective 6 IDS and traces		X		Ongoing	Section 6
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	

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 years 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). Y2 has seen several collaborations between the three research WPs. However, given the strong interdependence between the monitoring and configuration and repair operations, we focus here in reporting the collaborations between WP5 and WP6, as done in Y1. A detailed description about the currently ongoing collaborations and the recently ended ones can be found in section 3.2. Finally, please note that all fully integrated PhD students are listed in D8.2, including their co-supervisors and affiliation.

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 started last year was continued to integrate experienced researchers as well as new researchers not necessarily paid by FLAMINGO.

Table 5: Progress with respect to Y1.

Objective	Y1 activities	Y2 activities
S.M.A.R.T. Objective 1 Integration Ph.Ds	7 Ph.D.	14 Ph.D.
S.M.A.R.T. Objective 2 Research	37 papers	50 papers
WP Objective 1 Integration EU Research	NMRG; Dagstuhl AIMS; Coll. EU level	NMRG; Dagstuhl AIMS; Coll. EU level
WP Objective 2 Articles Online Systems	Not Addressed	Wikipedia contribution
WP Objective 3 Monitoring Architecture	Architecture for Security	Support for WP6
WP Objective 4 Flow Query Language	NFQL	Not Addressed
WP Objective 5 Monitoring Data	Internet Traffic Statistics SSH	Longitudinal analysis DNS; DNSSEC; SSH Link Dimensioning
WP Objective 6 IDS and traces	SSH attack tools fingerprinting	Compromise detection SSH Detection HTTP(S) detection
WP Objective 7 AI & Machine Learning	Not Addressed	HTTP adaptive streaming Botnet detection
WP Objective 8 ID & Fingerprinting	SSHCure	SSH Detection HTTP(S) detection
WP Objective 9 Cloud Infrastructures	DropBox analysis	Security and Clouds Cloud Networking

Table 7 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 of FLAMINGO and ended in Y2 because the research goals have been reached or they have branched into new collaborations. A collaboration is called **ONGOING** if started during Y1 or Y2 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.

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

Acronym	Researchers	WPs	Status
INRIA-JUB-RPL	A. Mayzaud - A. Sehgal	WP5, WP6	Ongoing
INRIA-JUB-Distr	A. Mayzaud - A. Sehgal	WP5, WP6	Ongoing
INRIA-UniBwM-Cloud	S. Seeber - G. Hurel A. Scherf	WP6	Ongoing
UCL-iMinds-Cache	D. Tuncer - M. Claeys	WP5, WP6	Ongoing
UT-UZH-Ethics	A. Sperotto - B. Stiller	WP7	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	Ongoing
iMinds-UPC-NetVirt	N. Bouten - R. Mijumbi M. Claeys	WP6	Ongoing
iMinds-UCL-Costs	B. Naudts - S. Verbrugge M. Charalambides - D. Tuncer	WP7	Ongoing
UCL-UT-MAN	D. Tuncer - R. Schmidt	WP5, WP6	Ongoing
JUB-UT-Pattern	N. Melnikov - R. Schmidt	WP5, WP6	Ended
UniBwM-JUB-RPL	S. Seeber - B. Stelte - A. Sehgal	WP6	Ended
UCL-UPC-BOSM	M. Charalambides - J. Rubio	WP7	Ended
iMinds-UT-QoS	R. Schmidt - N. Bouten	WP5, WP6	Ended
UZH-UPC-Policy	R. Garg - M. Flores	WP7	Ended
iMinds-UPC-Costs	B. Naudts - S. Verbrugge R. Mijumbi - M. Flores	WP7	Starting
iMinds-UT-ST	J. Famaey - A. Sperotto S. Latré	WP5, WP6	Starting
UT-UniBwM-Sec	TBD - S. Seeber	WP5, WP6	Starting
UT-JUB-Booters	J. Santanna - A. Sehgal	WP5	Starting
iMinds-UZH-VoS	B. Naudts - D. Doenni	WP7	Planned
INRIA-iMinds-WFlow	G. Hurel - N. Bouten	WP5, WP6	Planned
INRIA-UT-IoTSec	A. Mayzaud - A. Sperotto	WP5, WP6	Planned
UPC-UZH-Pricing	M. Flores - D. Doenni	WP7	Planned
UniBwM-UT-Social	D. Kergl - A. Sperotto	WP5	Planned
UPC-UniBwM-VLAN	R. Mijumbi - P. Hillmann	WP6	Planned
iMinds-UCL-SDN	S. Petrangeli - M. Charalambides D. Tuncer	WP5, WP6	Planned

3.2 Description of the Collaborations

This section presents the currently ongoing and recently ended collaborations between WP5 and WP6. Each collaboration description roughly follows the same structure. At first the topic of each collaboration is explained. Subsequently the progress and achievements in Y2 are highlighted. Depending on the status of a collaboration further steps are described. In the end each collaboration highlights the contribution to each WP.

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

Name	Affiliation	Collaborations	Acronym
Anthéa 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
Jair Santanna	UT	JUB	UT-JUB-Booters
Mario Golling	UniBwM	UT	UT-UniBwM-IDS
Sebastian Seeber	UniBwM	UZH, JUB, INRIA	UZH-UniBwM-JUB-M2 INRIA-UniBwM-Cloud
Achim Scherf	UniBwM	INRIA	INRIA-UniBwM-Cloud
Rashid Mijumbi	UPC	iMinds	iMinds-UPC-NetVirt
Anuj Sehgal	JUB	INRIA, UT, UniBwM	INRIA-JUB-RPL INRIA-JUB-Distr UT-JUB-Booters UZH-UniBwM-JUB-M2
Christos Tsiaras	UZH	UniBwM, JUB	UZH-UniBwM-JUB-M2
Daniel Dönni	UZH	UniBwM, JUB	UZH-UniBwM-JUB-M2
Niels Bouten	iMinds	UPC	iMinds-UPC-NetVirt
Maxim Claeys	iMinds	UCL	iMinds-UPC-NetVirt

bursts. (In the study, an extreme traffic burst is defined, at any timescale, by a traffic peak in which the traffic rate crosses a threshold given by the aggregate average rate plus three times the traffic rate standard deviation.) In addition, the results show that even in a more homogeneous network (i.e., hosts with similar access rates) one can identify extreme traffic bursts that might ultimately compromise Gaussianity fit. These bursts are typically caused by single, or very small group of hosts, and also by a small handful of applications (Applications studied by the collaboration were classified at the port level, e.g., HTTP in port 80 and HTTPS in port 443).

The collaboration contributes to WP5 since it is based on the collection of packet-level traffic traces from many different locations. This collaboration also contributes to WP6 since its results can be applied to management operations that directly or indirectly rely on Gaussian models.

3.2.2 Energy-aware Traffic Management (UCL-UT-Man)

Driven by the rising cost of energy and increasing environmental consciousness, some recent research efforts have been investigating the development of energy-saving resource optimization techniques for green operations. These aim at reducing the energy consumption of Internet Service Provider networks (and therefore operational expenditure) by adapting the configuration of network elements (i.e. line cards, interfaces or routers) according to traffic dynamics. While some efforts (e.g. [6]) have considered time-driven approaches by which network configuration profiles can be computed in a static fashion and applied according to the time period, others (e.g. [7]) have developed more adaptive approaches by which individual network devices can decide to enter sleep mode based on their awareness of current network conditions. Switching off entire links/routers, however, can adversely affect the network performance and subsequently cause degradation of

service quality. In particular, this can disconnect the network topology, which, in addition to possible packet losses under traffic variations, is not suitable for online reconfigurations given that the process of computing the routing configuration is not trivial.

In contrast to the majority of previous work that suggests switching off entire links, the adaptive energy-aware resource management approach developed [8] aims at reducing the energy consumption of core IP networks by applying the switching off strategy at the line card level. Exploiting the fact that many links in core networks are bundles of multiple physical cables, the proposed approach adapts the link capacity at run time by controlling the allocation of traffic in the network so that the load is distributed over a subset of router line cards, while unused ones enter sleep mode.

To achieve the energy-saving objective, the proposed approach relies on the DACoRM in-network management framework described in Deliverable D6.1 and presented in [9]. According to DACoRM, network edge nodes are embedded with a level of intelligence that enable them to realize self-management functionality and are organized into a management substrate through which they coordinate reconfiguration actions to control the traffic distribution in the network. More specifically, based on the path diversity provided by multi-topology routing (MTR), the network edge nodes adapt the volume of traffic routed over the multiple paths between any source-destination pair by adjusting the configuration of traffic splitting ratios. The main principle of the approach is to iteratively move the traffic load from less utilized line cards to more utilized ones that can accommodate this load and thus potentially fill-up their remaining capacity.

The performance of the energy management approach has been evaluated using real topologies and traffic traces, and the results show that substantial energy gain can be achieved without significantly compromising the balance of the network in terms of load [8]. Further evaluations have also been performed to investigate the influence of network topology characteristics (e.g. the number of nodes and degree of router connectivity) and MTR planes configuration on the performance of the adaptive scheme. In addition, the collaboration has investigated the influence of the routing protocol (e.g. multi-topology routing, Open Shortest Path First (OSPF), Equal-Cost Multi-Path (ECMP)) on network energy consumption.

The research issues related to the development of the resource management approach mainly fall within the scope of WP6.

3.2.3 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 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

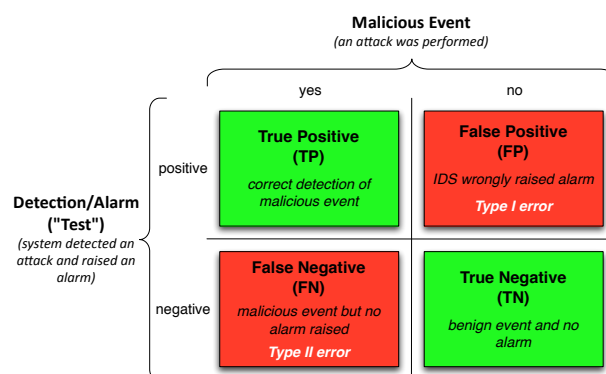


Figure 2: Categories of Alarms ("Confusion Matrix").

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].

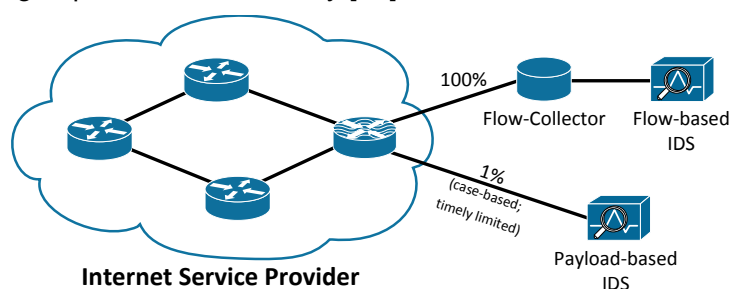


Figure 3: Simplified scenario for UT-UniBWM-IDS.

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 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,⁸ which detects SSH dictionary attacks and reports

⁸<http://github.com/sshcure/sshcure/>

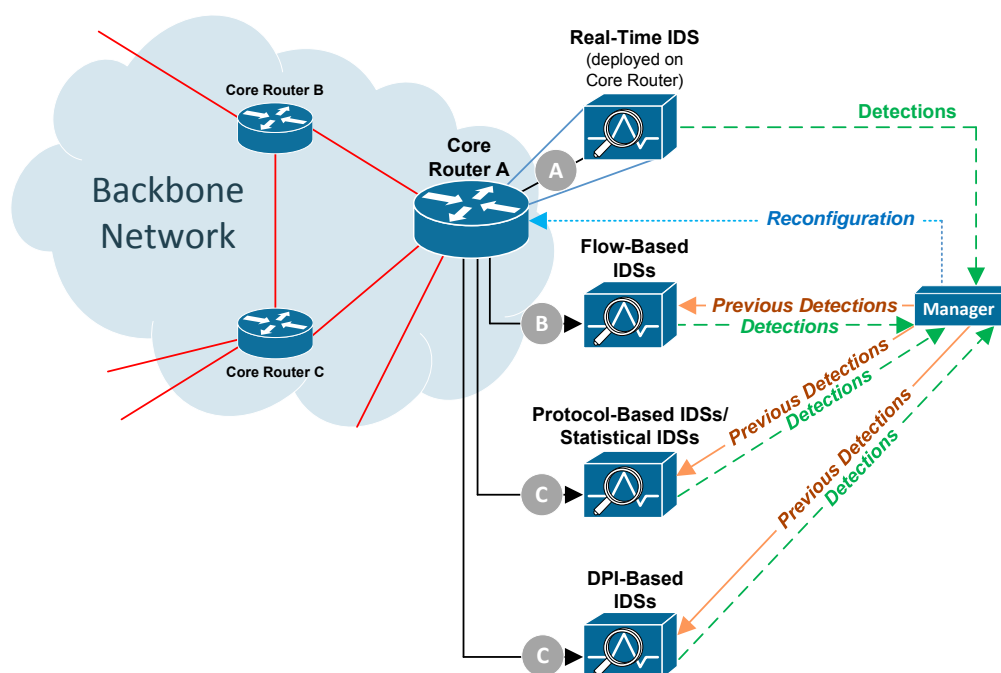


Figure 4: Architecture for multi-layered intrusion detection, from [13].

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.

For WP6, [16] also contributed to objective 3 "to develop an inventory of approaches for automated configuration and repair". By using cloud-based solutions (as described in [17, 18]), requirements for cloud-based services have been investigated and architectural approaches specific for this ap-

plication domain have already been partially developed (also addressing objective 9 “to propose and study automated configuration and repair in the context of the management of clouds (especially interclouds)”) As it is planned to develop an “automated” architecture that can be used in different administrative boundaries, objective 5 “to develop new architectures for automated configuration and repair approaches across administrative boundaries” is addressed as well.

3.2.4 Towards A Trust Computing Architecture for RPL in Cyber Physical Systems (UniBwM-JUB-RPL)

Cyber Physical Systems (CPSs) are widely expected to be formed of networked resource constrained devices. In order to suit the constraints of such networks, the Internet Engineering Task Force (IETF) developed the Routing Protocol for Low power and Lossy Networks (RPL) and Low-power and Lossy Networks (LLNs). Security in CPSs is important for maintaining the integrity and privacy of data, while also improving network resiliency to attacks. Even though RPL provides support for integrity and confidentiality of messages, details regarding key management and signatures are not covered. Since complexity and size is a core concern in LLNs, off-loading the security features to a Trusted Platform Module (TPM) makes it possible to include sophisticated security provisions in an RPL implementation.

This collaboration developed mechanisms to use the security mechanisms of a TPM in order to secure the communication in an RPL network. The design of a trust establishment and key exchange mechanism around the implied trust of a TPM to provide keys for secure RPL nodes, was a main task of this research. With this approach, the usage of a TPM on Resource Constrained Devices reduces the processing load on the main processor. The goal of this examination is the prevention of the dissemination of misleading routing information, which can affect the availability of the whole network.

The collaboration fits within WP6, since it develops a mechanism, specifically, targeted to RPL networks, to secure the communication. This approach is applicable in RPL networks which are used in wireless sensor networks.

3.2.5 QoE-Driven In-Network Optimization for Adaptive Video Streaming Based on Packet Sampling Measurements (iMinds-UT-QoS)

HTTP Adaptive Streaming (HAS) services allow the quality of streaming video to be automatically adapted by the client application in face of network and device dynamics. A major obstacle for deploying HAS in managed networks, is the purely client-driven design of current HAS approaches, which leads to excessive quality oscillations, globally suboptimal behavior, and the inability to enforce management policies.

These challenges can be tackled by steering the quality selection from within the network. iMinds already deployed a distributed in-network management heuristic, which is able to reduce the number of switches with a factor 5 and increase the quality up to 30%. One of the shortcomings of this deployment, however, was the assumption of a static bandwidth for each link and the absence of cross-traffic. To overcome this issue, this collaboration uses sampled packet measurements to measure and predict the per-link throughput, which is available for HAS traffic. Using these predictions the in-network video quality adaptation can divide the resources among the different HAS clients based on a provider’s policy.

This work contributes to WP5 because real world traffic traces will be collected and used to validate the link dimensioning approach.

This work contributes to WP6 in the following. The goal of this collaboration is to develop a distributed algorithm/heuristic, which is able to divide the resource among the various HAS clients subject to a providers policy. Using the measurements provided by WP5, each agent is able to perform a local optimization based on the throughput predictions. The different distributed agents share this network information and their local decisions with each other to be able to automatically react to changes in the network environment. Using the measurements and predictions on the current and future cross-traffic, the agents can make estimations on the residual bandwidth that can be shared among the different HAS sessions. These local estimations and the shared network information serve as input to the algorithm which limits the quality of the HAS sessions crossing the managed resources. This leads to a more stable quality selection at the client, since oscillations due to changed network environments are avoided.

3.2.6 Security of RPL Networks (INRIA-JUB-RPL)

The collaboration Security of RPL Networks between INRIA and JUB led to several sub-collaborations due to the existence of various types of attacks in RPL networks. In the following a summary of each individual topic is provided.

Mitigating DODAG inconsistency attacks is the first sub-collaboration. Fundamentally, RPL utilizes DODAGs, a directed graph like structure, to organize the routing topology in a network. The methodology used to detect and repair possible inconsistencies in DODAG can be manipulated by malicious nodes to harm the network.

The aim of this sub-collaboration is to develop methodologies to mitigate such attacks. An approach that dynamically adapts parameters of an adaptive threshold has been developed.

The second sub-collaboration is called **RPL Version number attacks**. Version numbers are used by the RPL DODAG root in order to keep track of the latest version of the topology. If a node detects that it is part of an older version, it is required to join the new version. However, due to the lack of security mechanisms, this method could be utilized by malicious nodes to attack the topology, and possibly even hijack nodes to join its own network.

The aim of this study is to evaluate the effectiveness of attacks based on manipulating version numbers, and also study the already proposed solutions. Based on the study a new approach that overcomes existing shortcomings would be developed.

Mitigating Black-hole and Sink-hole attacks is the last sub-collaboration. The goal is to develop a mechanism that mitigates black-hole and sink-hole attacks in RPL networks, by establishing inferred trust between neighbors.

This work is currently in the development phase, with network metrics contributing towards the trust metric already identified. An implementation of the preliminary approach is currently pending.

Monitoring of the RPL network and the identification of possible attacks in an RPL network contributes to WP5. The automated repair and mitigation of detected attacks in RPL networks contributes to WP6, as well as the application of developed approaches to wireless sensor networks, via the Internet of Things application area.

3.2.7 Distributed Monitoring Architecture for the Internet of Things (INRIA-JUB-Distr)

A generic distributed monitoring architecture is being designed for application in the Internet of Things (IoT) area. The goal of the architecture is to be able to monitor events and network flows passively without having any impact upon the resource constrained nodes that participate in such a network. This collaboration has been recently defined and is therefore in a preliminary stage.

The monitoring architecture will be applied to network and service monitoring in the IoT, and also towards anomaly detection and correction (including security aspects).

Developing a distributed monitoring architecture of the IoT infrastructure can be seen as a part of WP5. An automated repair of detected anomalies in the IoT and the application of developed approaches to wireless sensor networks, via the Internet of Things application area contributes mainly to WP6

3.2.8 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.

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.

3.2.9 Cache Management (UCL-iMinds-Cache)

The rise of Internet-based over the top multimedia services has put immense strain on the resources of Internet Service Provider (ISP) networks. This has resulted in increasing operating

costs but also in decreasing revenues of traditional traffic forwarding services. As a consequence, ISPs have started exploring alternative business models and service offerings. This has led to the deployment of Telco Content Delivery Networks (CDNs), which allow content to be cached deep inside the ISP network. For the operators, Telco CDNs reduce bandwidth demand on their backbone infrastructure and open up new business models. For the end-users, Quality of Service (QoS) can be significantly improved, as content is stored nearby and the ISP has full control over the network infrastructure. Furthermore, the advent of cloud computing and Software Defined Networking (SDN) technologies enable ISPs to virtualize their networks and by extension, their Telco CDN infrastructures. This opens up new business models and allows ISPs to dynamically offer virtual storage and content delivery services at the edge of the network, redeeming traditional CDN and content providers from installing additional hardware.

In this collaboration, iMinds and UCL have been investigating a cache management approach for multi-tenant caching infrastructures. The multi-tenant content placement and server selection problems were formally modelled by means of an Integer Linear Program (ILP) formulation. Solving this model determines where to store which content item of each tenant (content placement) and from which location to satisfy each request (server selection), with the objective of minimizing bandwidth consumption in the network while maximizing the cache hit ratio. The model can easily be adapted to consider other optimization criteria, such as cache hit ratio maximization or delivery delay minimization. All of these decisions are taken considering the content popularity and its geographical distribution. To compute a new configuration, the model requires predicted values concerning request distribution, for which a prediction strategy has been employed.

In the considered scenario, a large-scale ISP operates a limited capacity CDN service by deploying caching points within its network. Each network node is associated with caching capabilities, which enable a set of content items to be stored locally. The local caches can be external storage modules attached to routers or, with the advent of flash drive technology, integrated within routers. The ISP leases the caching space in its network to multiple content providers. Each content provider specifies the amount of caching capacity it wishes to lease for storing part of its content catalog.

The proposed approach has been thoroughly evaluated in a simulated environment on a Video-on-Demand (VoD) use case, for which a request trace of the VoD service of a leading European telecom operator was used. Its performance has been compared to a reactive caching approach, using the Least Recently Used (LRU) replacement strategy. The evaluations show that with the proposed ILP approach, the average total bandwidth usage inside the ISP network can be reduced by 12%-17%, depending on the amount of leased caching capacity. Furthermore, the average peak link usage can be improved by around 73%. Finally, evaluations showed that using the proactive ILP approach can lead to more balanced link load distribution inside the ISP network. All of these results are influenced by the accuracy of the prediction of content requests, which is the current focus of the collaboration.

Aspects related to the analysis of the VoD traces fall within the scope of WP5. This research activity mainly falls within the scope of WP6. The development of proactive mechanisms to efficiently manage the utilization of network resources concerns WP6.

3.2.10 Management of Virtualized Networks (iMinds-UPC-NetVirt)

The management of virtualized networks is a challenging task. Several technical challenges in terms of instantiation, operation, and management of virtual networks are either untouched or require further attention. Therefore, the collaboration is divided into sub-collaborations to focus on specific

research topics without losing the overall management goals existent in virtualized networks. In the following each sub-collaboration is explained:

Machine Learning-based Virtual Network Resource Management Most current approaches to resource management in network virtualization allocate a fixed amount of resources to the virtual nodes and links for their entire lifetime irrespective of actual utilization. As Internet traffic is not static, this could lead to an inefficient utilisation of overall network resources, especially if a substrate network rejects requests to embed new VNs while reserving the resources for VNs that are lightly loaded.

In this collaboration, instead of allocating a fixed amount of resources to a given VN throughout its lifetime, we dynamically and opportunistically allocate resources to virtual nodes and links depending on the perceived needs. We use a demand-driven dynamic approach that allocates resources to virtual nodes and links using machine learning techniques. Therefore, after the initial virtual network embedding step, resources allocated to each virtual node and link are monitored and adjusted to reflect both actual resource need by the virtual network, and resource availability in the substrate network. We represent each substrate node or link as an agent. These agents are tasked to monitor the resource utilisation of all mapped virtual nodes and links, and comparing this with the available substrate resources, re-allocations are performed. The agents then monitor the network to determine its performance, for example through parameters such as packet delay and drop ratio, and based on these statistics, the agents use machine learning techniques to make better actions for future resource allocations. We have been able to apply three machine learning techniques to this problem: Reinforcement Learning [24], [25],[26], Artificial Neural Networks [27], [26] and Neuro-fuzzy Algorithms [28]

SDN-based Management of Virtual Network Resources In this collaboration, we use the SDN control plane to efficiently manage resources in virtualized networks by dynamically adjusting the virtual network (VN) to substrate network (SN) mappings based on network status. We extend an SDN controller to monitor the resource utilisation of VNs, as well as the average loading of SN links and switches, and use this information to proactively add or remove flow rules from the switches. The results are presented in [29].

Chaining of Service Functions in Network Function Virtualisation In this part of the collaboration the different network functions involved in multimedia services are identified and how they are chained to support the delivery of multimedia services. The different opportunities for virtualizing multimedia network functions will be determined and how this impacts the service function chains. We will determine which physical resource deployments are optimal for such virtual multimedia services (i.e. locations of datacenters, peering points). Furthermore, the optimal network function mappings will be determined based on the alternative service function chains and the resource availability at the substrate network. This work is still in progress and has not yet resulted into a submitted/accepted paper.

The overall research activity of this collaboration falls within the scope of WP6, since automated and dynamic reconfiguration of virtualized networks based on utilization information is a main goal.

3.2.11 Cloud Security (INRIA-UniBwM-Cloud)

The aim of the joint research activity Cloud Security between INRIA and UniBwM is to investigate recently available SDN-based mechanisms for delivering security in different network scales, ranging from home networks to datacenters. In the first step the scope of the study is focused on several well-known network attacks such as denial-of-service, information gathering and malware propagation.

While mainly the recent Openflow-based solutions for mitigating such kind of attacks are emphasized, the collaboration also explores some previous SDN attempts - such as ForCES and Active Networks - in this light. Furthermore, the proposed approaches will be compared with the ones that are found nowadays in traditional networks (i.e. non-SDN). In addition, several well-known Openflow controllers are evaluated to identify the most suited ones for implementing security solutions in SDN networks. As a result a framework to compare OpenFlow compatible SDN controllers regarding their security functions will be presented.

The active collection of network flow metrics and OpenFlow statistics for security monitoring builds the basis for this collaboration and therefore concerns WP5. Related to WP6 an automated mitigation and network (i.e. forwarding-plane) reconfiguration is done once a network attack has been detected.

3.2.12 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 [30].

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/NfSen⁹ 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 [31],¹⁰ which adds a geographical dimension to flow data and displays the data on a map. Since SURFmap [31] 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.

Several technical aspects of this collaboration have been done in Y1. The main activity carried out in Y2 is related to improving the Internet Draft (ID) ([30]) based on community feedback. 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.

⁹<http://nfsen.sourceforge.net>

¹⁰<http://surfmap.sourceforge.net>

4 Data Collection

Among the FLAMINGO partners, several data collection activities are taking place. In this section, we provide an overview of the type of traces that are being collected in the context of the Joint Security Lab (Section 4.1). Then we present several datasets that have been publicly released as a result of research activities (Section 4.2).

4.1 Data Collection and Joint Security Lab

Several data collection activities are linked to the Joint Security Lab. Although such measurements are typically not made public, they can be shared with the consortium partners upon request and when a researcher is visiting an institution part of the Joint Security Lab. For more information regarding the Joint Security Lab, we refer the reader to D1.2. In the following (Table 8) we report the main data collection activities carried on by the FLAMINGO partners. The highlighted cells refer to measurements that have been started in Y2, while the other measurements were already active in Y1 and have continued in Y2.

Table 8: Traffic monitoring activities within the Joint Security Lab.

Traces	UT	INRIA	UniBwM	iMinds	JUB	UZH
NetFlow/IPFIX	X	X	X			
sFlow	X		X			
SNMP	X					
Honeypots data	X	X	X			
DNS	X					
IDS alerts		X	X			
Blacklists	X	X				
Network Telescope			X			
Server Logfiles	X		X			
Social Networks			X			
Video on Demand				X		
Mobile QoS			X (public)		X (public)	X (public)
Ripe Atlas (traceroute)	X (public)					

4.2 Publicly-released Datasets

Some research activities in WP5 have resulted in the public release of datasets. This is the case, for example, for the paper [32] on the possible misuse of DNSSEC for DDoS amplification attacks (published at IMC 2014) and the paper [2] on SSH flow-based compromised detection (published in ACM CCR). In addition, public data has been made available for the research on mobile QoS by the Bonafide application,¹¹ developed in collaboration between JUB, UZH and UniBwM. In this section, we summarize these datasets.

- **DNSSEC and its Potential for DDoS Attacks** – UT and SURFnet Bv have conducted a large-scale measurement study apt to gauge the possible impact of DNSSEC signed domains for use in DDoS reflection and amplification attacks. The study has made use of active

¹¹<https://bonafide.pw>

measurements based on DNS zone files for the top-level domains .com, .net, .org, .nl, .se and .uk, effectively covering 70% of all 3.5 million DNSSEC-signed domains. The dataset is released in anonymized form and it is available for download at <https://traces.simpleweb.org/>. For more information about this research, please see Section 8.

- **SSH Dataset for Intrusion Detection** – The SSH datasets feature a unique combination of flow data (exported using NetFlow) and authentication log files, allowing for validation of flow-based intrusion detection systems. The dataset consists of two groups of traces, each one covering one month of flow data and including the respective log files for authentication validation. The flow data has been exported by a Cisco Catalyst 6500 with SUP2T supervisor module (PFC4, MSFC5), and collected using `nfcapd`.¹² Neither packet sampling nor flow sampling have been applied, and only SSH traffic has been included. The log files have been gathered from various Linux, BSD and Mac OS X operating systems. The dataset is available in anonymized form for download at <https://traces.simpleweb.org/>. More information about the research related to this dataset can be found in Section 6.
- **Bonafide+ mobile QoS** – The Bonafide+ project is targeting QoS measurement collection in mobile networks. Data is collected via an Android application that performs protocol-specific and random data flow-based measurements against geographically distributed endpoints. All information collected by the BonaFide+ application, once anonymized, can be downloaded by researchers to process on their own. The BonaFide+ application uses the collected information to derive Quality of Experience (QoE) scores, however, the recorded information includes all raw results, such as latency, bandwidth, signal strength and location for every test. The recorded information also includes results from randomized data flows, which are unlikely to be susceptible to traffic shaping, and also protocol specific results. Network operator-specific information is also recorded in the results. The collected measurements are available at <https://bonafide.pw/web/results.html>.

¹²<http://nfdump.sourceforge.net/>

5 A Longitudinal Analysis of Internet Rate Limitations

In this section and in the following Sections 6–9, we summarized selected highlights of the research conducted in WP5 during Y2.

TCP is the dominant transport protocol for Internet traffic, but the preponderance of its congestion control mechanisms in determining flow throughput is often disputed.¹³ In the context of the Flamingo project we have performed a study to analyze the extent to which network, host and application settings define flow throughput over time and across autonomous systems. Drawing from a longitudinal study spanning five years of passive traces collected from a single transit link, our results show that continuing OS upgrades have reduced the influence of host limitations owing both to windowscale deployment, which by 2011 covered 80% of inbound traffic, and increased socket buffer sizes. On the other hand, we show that for this data set, approximately half of all inbound traffic remains throttled by constraints beyond network capacity, challenging the traditional model of congestion control in TCP traffic as governed primarily by loss and delay.

5.1 Dataset and Preprocessing

The analysis was conducted using unanonymized traffic traces of the MAWI dataset [34], which groups traffic traces captured at the Japanese National Research and Education Network WIDE. The dataset consists of daily 15-minutes traffic traces. For the scope of this research, traces spanning the 5-year period 2006-2011 have been analyzed.

Table 9: Overview of traced MAWI dataset.

Year	Days	TCP data flows ($\times 10^6$)	Traffic In (TB)	Traffic Out (TB)	ASes ($\times 10^3$)	Prefixes ($\times 10^3$)
2006	91	20.52	0.43	0.45	10.90	56.86
2007	350	102.56	2.11	2.49	17.21	113.79
2008	358	112.26	2.43	2.10	24.74	156.54
2009	364	113.97	2.48	2.53	19.71	143.87
2010	365	113.70	2.58	3.43	20.38	148.03
2011	358	114.74	3.44	5.14	19.99	140.56
Total	1886	5777.55	13.50	16.14	34.12	341.22

An overview of the dataset used is provided in Table 9. It comprises 5.7 billion flows over five largely uninterrupted years, which represents approximately 30 terabytes of TCP traffic. For the purposes of this work the focus is exclusively on inbound traffic (i.e. traffic towards a destination in the WIDE network), 60% to 80% of which originates from port 80.

All TCP flows are reassembled and analysed for each daily trace. In addition to the five tuple used to define each connection, we impose two additional restrictions: a contiguous sequence number space and a three minute timeout. These restrictions are helpful to deal with port reuse and unterminated flows respectively. For the reconstruction of TCP flows we have used a custom tracer since the MAWI traces impose two constraints which require careful consideration: the proportion of bidirectional flows (40%-60% over 5 years) and the short duration of each individual trace file.

¹³The research summarized in this section has been published in the paper [33]. Araujo, R. Landa, R. Clegg, G. Pavlou, and K. Fukuda *A longitudinal analysis of Internet rate limitations* In Proceedings of IEEE International Conference on Computer Communications (INFOCOM 2014)

Loss is inferred by accounting for retransmissions in the upstream data and out-of-order packets in downstream data; in the remainder of the text we will refer to the end-to-end loss as the sum of out-of-order and retransmitted data bytes over the total data bytes in a given direction.

Each daily trace in the dataset is processed from a packet level capture into a collection of flow level statistics. This gives us insight into the end-to-end characteristics of traffic. However, since a core objective of this work is to augment this time-based information with data describing the endpoints of each flow, aggregating by location is also required. Location information is added by mapping the original source and destination IP addresses to its geographical and topological counterparts. We use the *routeviews* archives to reconstruct the mapping between each IP and both AS and network prefix. Mapping IP to country is done through the use of GeoLite, a commercial geolocation database. After associating flows to country, region, AS and network prefix for both source and destination IPs, we aggregate flow statistics over each location identifier. This generates a daily collection of location identifiers and associated flow properties, from which we can sketch the geographic and topological properties of the dataset over time.

5.2 Flow Classification

One fundamental precondition to decouple the influence that network loss, host configuration and TCP behaviour has on the throughput experienced by a flow is the reconstruction of the congestion window behaviour of TCP flows on the basis of observed data. Unfortunately, the congestion window value is internal to the senders TCP state machine and may not manifest itself in the absence of sufficient data from the application layer. A more easily observed quantity which serves as a reasonable proxy for the congestion window is the number of unacknowledged bytes in flight, henceforth referred to as the flight size, which can be derived given an accurate estimate of the end-to-end delay. The evolution of both flight size and RTT can in turn be used to ascertain to what extent throughput is regulated by limitations imposed at different layers of the networking stack. Having obtained flight information from each flow (method can be found in [33]), we next consider what is the predominant factor that affects its throughput. Within the context of TCP, we classify flows as being artificially constrained by three distinct processes: *application pacing*, *host limited* and *receiver shaping*.

Application Paced Flows - we identify flows as being application paced if the period between bursts terminated by application limited flights is consistently under 10 seconds and the standard deviation of the intermediate pauses is under one second. This definition is purposely designed to reject flows which exhibit long silence periods due to user interaction, and follows closely the behaviour historically associated with Youtube video streaming.

Host Limited Flows - given sufficient bandwidth and traffic to send, a flow may encounter local constraints at either end-host which caps its throughput. For instance, the buffer space allocated on both the sender and receiver side is often pre-configured, and it is common practice to tune these values down on popular servers and managed infrastructure in a bid to conserve memory or bandwidth. Host limited cases are characterised by a constant window size over time.

Receiver Shaped Flows - a flow which is neither application paced or host limited can still be artificially constrained by flow control. In addition to the sender, the receiver can also shape throughput by manipulating the advertised window announced on every acknowledgement. We classify flights as being receiver-shaped if the cross- correlation between the advertised window size and the maximum flight size is statistically significant with a p-value less than 0.05.

5.3 Highlights of Analysis Results

After processing each daily trace individually, the results were aggregated longitudinally in order to trace the evolution of constraints affecting TCP across both time and spatial/topological dimensions. The analysis is framed as a re-visiting of four commonly held assumptions regarding Internet throughput. The aim is to provide much a needed factual verification of these assumptions, which itself can lead to a re-appraisal of Internet throughput modelling efforts. This section provides an overview of our main findings from the analysis; more elaborate analysis can be found in [33].

Table 10: Percentage of traffic bytes affected by each constraint by year.

Year	Application (%)	Host (%)	Receiver (%)	Total (%)
2007	49.47	18.58	0.55	68.60
2008	49.55	17.80	0.69	68.04
2009	47.10	14.50	2.57	64.17
2010	36.78	20.44	3.21	60.43
2011	46.10	13.49	0.60	60.20

Table 10 displays the extent to which each of the three limitations affects inbound traffic in the MAWI dataset over time. The bulk of the volume in bytes is either conditioned by application pacing or host limits - the former tends to be the largest type of limitation because it affects streaming traffic. The use of receiver shaping on the other hand is both small in scale and temporally confined to 2009 and 2010. Over five years, the overall effect of the three selected constraints has dropped by close to 10%. Below, we compare our findings against commonly held beliefs.

Assumption A. Throughput is primarily shaped by TCP

Overall, we found that flow rates are not typically dictated by TCP congestion control alone. While the impact of application pacing and host limitations on rate control is decreasing, as of early 2012 we found that less than 40% of all inbound traffic had TCP congestion control as the primary rate control mechanism. The reduction of application pacing however may reflect the nature of the observed link, as many traditional streaming providers have migrated towards peering or CDNs, bypassing interdomain links entirely. As such we expect the effect of application pacing to be more pronounced when considering traffic beyond transit.

Assumption B. Throughput is primarily sender driven

The endpoint which ends up dictating the maximum achievable throughput through flow control is typically a function of the OS adoption cycle. With the window scale option covering 80% of all inbound traffic, the main source of host level constraints are now conservative buffer sizes. For this dataset, hosts internal to WIDE have seemingly been upgraded at a faster rate, or less conservatively, than their remote counterpoints. As such, throughput has become increasingly sender driven over time for inbound traffic.

Assumption C. Throughput is correlated with flow size

The analysis results confirm the notion that the highest throughputs are attained by the largest flows, but they also show that improvements in throughput do not apply equally to all flow sizes. Whereas throughput has consistently improved for low-volume traffic, it has not done so for high-volume traffic. Hence, these findings suggest an increased differentiation between high-value, low-volume traffic whose throughput has markedly increased, and low-value, high-volume traffic whose throughput has stagnated.

Assumption D. Throttling primarily affects heavy hitters

Receiver shaping was limited in both scope, affecting at most 4% of bytes, and time, being primarily concentrated within 2009 and 2010. Prior to 2009, receiver shaping mostly targeted flows which attained the highest throughput, and may have in part been performed by hosts. By 2011, receiver shaping mostly subsided; the selected targets of shaping were neither the biggest contributors in terms of volume, nor the most aggressive senders.

6 SSH Compromise Detection using NetFlow/IPFIX

Brute-force attacks against SSH daemons are a common type of attack in which attackers perform authentication attempts on a remote machine.¹⁴ These attacks are often performed using dictionaries – lists of frequently used login credentials – and are therefore commonly referred to as *dictionary attacks*. Once compromised, machines can cause serious damage by joining botnets, distributing illegal content, participating in DDoS attacks, mining Bitcoins or other cryptocurrencies, etc. The threat of SSH attacks was recently stressed again by the *Ponemon 2014 SSH Security Vulnerability Report*: 51% of the surveyed companies has been compromised via SSH in the last 24 months [35]. From our own experience, we know that compromises can have various causes, among which are dictionary attacks. For example, in campus networks, such as the network of the University of Twente (UT) with roughly 25k active hosts, we observe approximately 115 dictionary attacks per day, while in a backbone network, such as the Czech National Research and Education Network CESNET, it is not uncommon to observe more than 700 per day. Even more attacks should be expected in the future; several renowned organizations, such as OpenBL¹⁵ and DShield,¹⁶ report a tripled number of SSH attacks between August 2013 and April 2014. Research in the area of compromise detection is therefore crucial to reduce the potential damage caused by compromises.

Flow-based approaches for SSH intrusion detection have been developed to overcome the scalability issues of host-based alternatives. Although the detection of many SSH attacks in a flow-based fashion is fairly straightforward, no insight is typically provided in whether an attack was successful. We address this shortcoming by presenting a detection algorithm for the flow-based detection of compromises, i.e., hosts that have been compromised during an attack.

This research is conducted by performing a structured analysis of, on the one side, the flow-based behavior of the most-commonly deployed SSH demons and host-based SSH mitigation mechanisms, and on the other side, the flow-based behavior of several SSH dictionary attack tools. In the following, we introduce the flow-based SSH attack model on which this research is based and the proposed *compromise detection* algorithm, followed by the validation results.

6.1 The detection model



Figure 5: SSH attack phase transitions.

The work on the detection algorithm and our IDS *SSHCure*¹⁷ has started in 2011, when we developed a detection algorithm based on the Hidden Markov Model (HMM) of SSH attacks presented in [36]. That model assumed SSH attacks to feature one or more of the following three attack phases:

- **Scan** – An attacker scans hosts for active daemons.

¹⁴The research summarized in this section has been published in the paper [2] R. Hofstede, L. Hendriks, A. Sperotto and A. Pras *SSH Compromise Detection using NetFlow/IPFIX*, ACM CCR 2014

¹⁵<http://www.openbl.org>

¹⁶<http://www.dshield.org>

¹⁷<https://github.com/SSHCure/SSHCure>

- **Brute-force** – An attacker performs many authentication attempts on target hosts, using a large number of username/password combinations (dictionary).
- **Compromise**¹⁸ – An attacker has gained access to a target host by using correct login credentials.

The SSH attack state transitions are shown in Figure 5. The only difference with the work in [36] is that attacks can also start in the *brute-force* phase for the following reasons. First, the *scan* phase can have taken place in the past (e.g., before traffic observation has started). Second, our investigation of attack tools has shown that attacks can start directly from the *brute-force* phase, as the scan can have been performed by another host or the target may be known in advance by an attacker.

It was also shown in [36] that the attack phases can be clearly observed in a time-series of the number of packets-per-flow (PPF). The *scan* phase is characterized by a low number of PPF, indicative for TCP connections that do not complete their handshake. At some point in time, the attacker starts the *brute-force* phase, trying to authenticate to remote machines. This phase was shown to consist of flows with a significantly larger number of PPF, caused by the authentication procedure of SSH. After a potentially successful login, some residual traffic can be observed between attacker and target as part of the *compromise* phase, where traffic is characterized by having a PPF that is outside the range of what is considered *brute-force* phase traffic.

A purely PPF-based detection algorithm and a first prototype of *SSHCure* were presented in [37]. Although this yielded promising results, a large number of production deployments of *SSHCure* has proven that the designed detection algorithm has various shortcomings that ultimately cause compromises to remain undetected or false alarms to be raised.

6.2 The detection algorithm

The combined analysis of the SSH demons, SSH attack tools and mitigation mechanisms has brought to the definition of a set of scenarios describing the flow-based behavior of a successful compromise and to the definition of a compromise detection algorithm.

The analysis of successful compromises has shown that SSH attack tools can all be described based on four actions that happens upon compromise. These actions are as follows:

- *Maintain connection, continue dictionary* – The connection with successful authentication is maintained, until the end of the attack. The attacker continues with the attack, also towards the compromised host.
- *Maintain connection, abort dictionary* – The connection with successful authentication is maintained, until the end of the attack. Other attack traffic towards the compromised host is stopped.
- *Instant logout, continue dictionary* – The connection with successful authentication is closed right after the compromise, while the attacker continues to attack both the compromised host and others.
- *Instant logout, abort dictionary* – The connection with successful authentication is closed right after the compromise. Other attack traffic towards the compromised host is stopped.

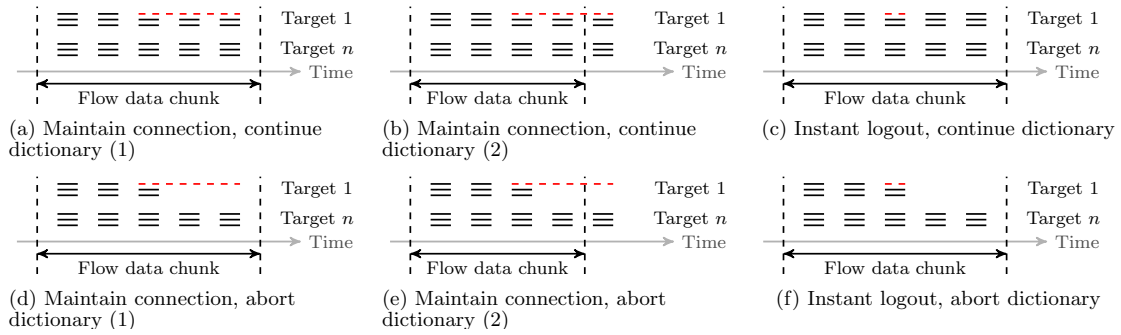


Figure 6: Various types of compromise flows in a chunk of flow data.

Key to our compromise detection are the four actions that can be observed after a compromise. We have transformed these actions into six scenarios, as shown in Figure 6. The two additional scenarios have been defined to accommodate for the fact that many analysis applications receive and process flow data in fixed-size time bins, as a consequence of which our algorithm has to take into account that attack data may be spread over multiple data chunks. Each of the subfigures shows a flow data chunk, with flows (long dashes) towards targets running an SSH daemon. Short-dashed lines mark a flow with a compromise. The compromise detection algorithm that we propose identifies a possible compromise if a traffic pattern is recognized that matches one of the aforementioned scenarios.

6.3 Validation results

The validation of the compromise detection mechanisms has been carried on based on two one-month long datasets. Each dataset consists of both network flows and server-side logs (see also Section 4.2). The first dataset has been collected by monitoring a set of low-, medium- and high-interaction honeypots. The second dataset comprises instead the traffic and logs of around 100 servers at the University of Twente campus.

Table 11: Validation results per dataset

	TPR	TNR	FPR	FNR	Acc
D1 (honeypots)	0.692	0.921	0.079	0.308	0.839
D2 (servers)	–	0.997	0.003	–	0.997

For each of the datasets we report the performance metrics of true/false positive (TP and FP) rates and true/false negative (TN and FN) rates, and the overall detection accuracy, defined as $Acc = \frac{TP+TN}{TP+TN+FP+FN}$. The overall detection results are summarized in Table 11. The detection results highlight that the achieved accuracy for D1 is 83%. The reason for this result is to be found in the specific honeypot setup. Many of the honeypots are hosting specially crafted SSH daemons with easily guessable login and password, with the consequence that several SSH brute force attacks will record a successful login at the first attempt. In this case, the compromise will look the same of a legitimate authentication. D2 has instead an accuracy close to 100%.

¹⁸Although we use the attack phases presented in [36], we use the more intuitive *compromise* to denote the *die-off* phase.

7 DNSSEC Meets Real World: Dealing with Unreachability Caused by Fragmentation

The Domain Name System (DNS) provides a critical service on the Internet: translating host names into IP addresses.¹⁹ Traditional DNS does not provide guarantees about authenticity and origin integrity. DNSSEC, an extension to DNS, improves this by using cryptographic signatures, at the expense of larger response messages. Some of these larger response messages experience fragmentation, and may, as a result of that, be blocked by firewalls. As a consequence, resolvers behind such firewalls will no longer receive complete responses from name servers, leading to certain Internet zones becoming unreachable because no translation into IP addresses can be performed.

This research shows that despite ongoing efforts to educate firewall and resolver administrators, as much as 10% of all resolvers suffer from fragmentation-related connectivity issues. Given that some major Internet companies were reluctant to adopt even a technology like IPv6 if it meant that a small percentage of their users would have connectivity issues, it is clear that we cannot rely on resolver/firewall operators alone to tackle this issue.

The contribution of the paper [1] is that it a) quantifies the severity of these DNSSEC deployment problems, based on extensive measurements at a major National Research and Education Network (NREN) and backed up by validation of these findings at an independent second location, b) proposes two potential solutions at the DNS authoritative name server side, and c) validates both solutions, again based on extensive measurements on the operational network of this major NREN. The paper concludes with a recommendation favoring our first solution. The first solution is relatively simple to implement and gives DNS zone operators control over this problem without having to rely on all resolver operators solving the issue.

7.1 DNSSEC and Fragmentation

We verify the presence of fragmentation and quantify the extension of the problem by means of traces collected on an authoritative name server of SURFnet. This server is authoritative for ± 4000 zones, including ± 300 DNSSEC-signed zones. It receives ± 500 queries per second on average. The traces recorded in early 2012 over a period of 6 hours contain about 8.5 million DNS(SEC) messages.

Problematic Resolver Characteristic	Occurrence
CASE 1: Send ICMP Fragment Reassembly Time Exceeded	1.3%
CASE 2: Fallback to traditional DNS	2.4%
CASE 3: Reduce advertised max. response size in retries	3.5%
CASE 4: TCP fallback w/o truncated UDP response	<0.1%
CASE 5: Retries for large responses (>512 bytes)	9.7%

Table 12: problematic resolver characteristics.

The research in [1] identifies a set of five behavioral patterns indicating that a resolver is experiencing fragmentation problems (problematic resolver). Table 12 summarises these behavioral patterns and quantifies the occurrences of each of those in the collected traces. In overall, the trace analysis

¹⁹The research summarized in this section has been published in the paper [1] G. van den Broek, R. van Rijswijk-Deij, A. Sperotto, and A. Pras. *DNSSEC meets real world: dealing with unreachability caused by fragmentation*. IEEE communications magazine, 52(4):154160, April 2014

showed that as much as 10.5% of the resolvers in the trace experience problems related to fragmentation. The same analysis has also been conducted on an authoritative name server at the University of Pennsylvania and showed the same distribution of behaviours.

7.2 Solutions to Fragmentation

In the following, we discuss two possible solutions to the fragmentation problems and their respective strengths and drawbacks.

7.2.1 Avoiding Fragmentation in General

A simple solution to the problem is to avoid fragmentation in general. This can be achieved by acting on the advertised maximum response size for EDNS0 (the extension to DNS that allows responses larger 512 bytes [38]). Problematic resolvers advertise maximum response sizes that are too high (as they cannot receive fragmented responses). The solution we propose here is to restrict the maximum response size in the configuration of the authoritative name server, such that (most) response fragmentation is avoided. When responding to queries the authoritative name server will then use the minimum of the configured response size and that advertised in the query. The response size limit will help most problematic resolvers, even if just one of the authoritative name servers per zone returns responses of limited size. This is because resolvers query all authoritative name servers for a zone in case they do not receive a response. The advantage of this solution is its simplicity, although care should be taken in choosing an appropriate response size. The downside of this solution is that it affects all resolvers independently if they show a problematic behaviour.

7.2.2 Selectively Avoiding Response Fragmentation

A different solution is avoiding response fragmentation by limiting the response size for problematic resolvers only. This solution is based on DNSRM7 (DNS Router/Modifier), a tool we developed specifically for this purpose. DNSRM operates as a host-proxy on an authoritative name server and acts on information supplied by a separate sensor tool that detects problematic resolvers based on a set of behavioral patterns. The purpose of DNSRM is to allow an authoritative name server to differentiate in response size, depending on the querying resolver. The advantage of this solution is that it affects problematic resolvers only. However, it is also more invasive since it requires DNSRM to run on an authoritative name server.

7.3 Validation Results

Both solutions were tested for 6.5 hours during office hours on the same authoritative name server of SURFnet that had been used for the preliminary analysis in Section 7.1. The conducted validation showed the following:

- When fragmentation is avoided in general, the traces show that no responses were fragmented at the authoritative name server. Consequently, we saw no ICMP FRTE messages, which indicates that this solution effectively helped the hosts that were previously sending these error messages.

- When fragmentation is avoided in a selective manner, the traces show that fragmentation was down about 50% compared to normal operations; note that fragmented responses still occur for resolvers that are not marked as problematic resolvers. Also, the number of ICMP FRTE messages was 18% of normal. This number is not zero, because the sensor first needs to detect a problematic resolver. Only after detection will the resolver be helped and will these ICMP messages disappear.

Although both solutions proved to be effective, the overall evaluation indicates that the first solutions should be preferred in operational environments due to its simplicity and effectiveness. For more details about this research, we refer the reader to [1]

8 DNSSEC and Its Potential for DDoS attacks

Over the past five years we have witnessed the introduction of DNSSEC, a security extension to the DNS that relies on digital signatures. DNSSEC strengthens DNS by preventing attacks such as cache poisoning. However, a common argument against the deployment of DNSSEC is its potential for abuse in Distributed Denial of Service (DDoS) attacks, in particular reflection and amplification attacks.²⁰ DNS responses for a DNSSEC-signed domain are typically larger than those for an unsigned domain, thus, it may seem that DNSSEC could actually worsen the problem of DNS-based DDoS attacks. The potential for abuse in DNSSEC-signed domains has, however, never been assessed on a large scale. The goal of this research is to establish ground truth around this open question.

8.1 DNS Amplification

Fig. 7 shows how DNS reflection and amplification attacks work. Reflection attacks leverage on *spoofing* to be able to re-direct unwanted DNS responses to the victim target. Attacks are initiated from a swarm of machines (left-hand side of the figure) under the control of the attacker. The attacker uses this swarm to send large numbers of DNS queries in which the sender IP address is spoofed to be the victim's IP address (bottom middle of the figure). Queries are sent ① to so-called *open DNS resolvers*. These are misconfigured DNS resolvers that do not restrict which clients are allowed to send them queries. In turn, the open resolvers will – if the query result is not cached – contact the appropriate authoritative DNS servers ② to resolve the query. Finally, the open resolvers will send the responses ③ to the victim. In general, the queries ① are small whereas the responses ③ are large, hence achieving amplification. Amplification is defined as $\frac{\text{response size}}{\text{query size}}$. Typical DNS requests are in the order of magnitude of 20 – 60 bytes in size. The classic DNS protocol [39] limits responses to at most 512 bytes; assuming a request size of 40 bytes this already yields an amplification factor of $\frac{512}{40} \approx 12.8$. More recent extensions to DNS that allow for larger responses easily result in amplification factors of 100 or more.

Unfortunately, open resolvers are plentiful on the Internet. Kühner et al. [40] report observing between 23 and 25.5 million open resolvers during weekly Internet-wide scans over a period of 4 months between November 2013 and February 2014. This makes this attack easy to carry out and thus attractive for attackers.

8.2 Measurement Methodology

Table 13 summarized the query performed for each signed and unsigned domain. The responses to the different query types will have different sizes which in turn will lead to different amplification factors. For the three query types A, AAAA and TXT, we perform multiple queries. One for the so-called apex record (denoted by @), one for the name `www` under the domain and one for the name `mail`. We expect that at least one of these names exists in most domains. Each individual query is performed once using classic DNS and once using EDNS0 with the EDNS0 maximum response size set to 32768²¹. For DNSSEC-signed domains we also perform the query with EDNS0 and the DNSSEC OK (DO) flag set to `true`, to get DNSSEC-signed responses. For each queries, several

²⁰The research summarized in this section has been published in the paper [32] R. van Rijswijk-Deij, A. Sperotto, A. Pras *DNSSEC and its potential for DDoS attacks*, ACM IMC 2014

²¹We chose this value to also register results that exceed the commonly used maximum response size of 4KB; we decided not to use the maximum value (65535) since we did not want to risk running into possible boundary conditions in DNS software implementations.

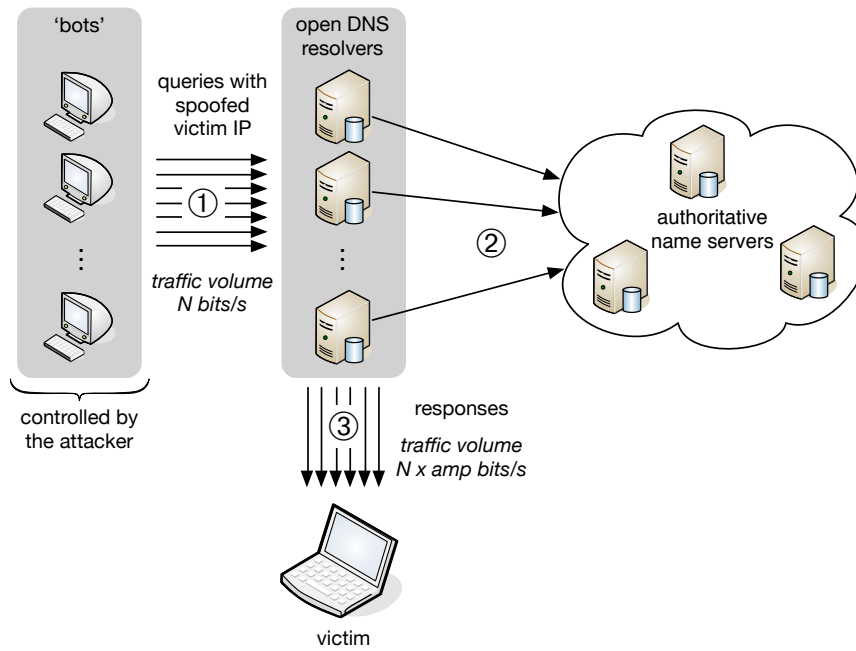


Figure 7: DNS amplification attack

Signed and unsigned domains	
Query type	Description
ANY	the query includes all resource records for the queried name
MX	the query returns the names of mail exchangers for the domain
NS	the query returns all authoritative name servers for a domain
A	the query returns the IPv4 address(es) for the queried name
AAAA	the query returns the IPv6 address(es) for the queried name
TXT	the query returns textual information for the queried name
Signed domains only	
Query type	Description
DNSKEY	the query returns the set of public keys required to validate signatures in a domain.
NSEC(3)	(<i>authenticated denial of existence</i>) this records proves with a digital signature that the queried name does not exist.

Table 13: Performed query for each selected domain

metrics are recorded. For a complete description of those, we refer the reader to [32]. In the following, we focus on the amplification factor.

8.3 Data Sets

We summarize in the following the source data retrieved from the ToD and the data that resulted from our experiments.

TLD	Data obtained	#domains	#DNSSEC
.com	Full zone	113.1M	326.5k (0.3%)
.net	Full zone	15.2M	69.5k (0.5%)
.org	Full zone	10.3M	37.6k (0.4%)
.nl	Selection	5.4M	1696.1k (31.2%)
.se	Full zone	1.4M	334.9k (24.8%)
.uk	Selection	10.6M	10.2k (0.1%)

Table 14: Overview of source data

8.3.1 Source Data

We obtained data covering 70% of all 3.5 million DNSSEC-signed domains from six different top-level domains. Tab. 14 lists the TLDs from which we obtained data. The table lists the type of data obtained (either the full zone or a selection containing all secure delegations and a random sample of unsigned domains), the total number of domains in the TLD and the number of secure delegations (as an absolute value and a percentage).

8.3.2 Collected data

TLD	#domains	#failed	#skipped	#queried	#queries	#auth ns
.com	326504	7416	471	318576	54.6 M	2550
.net	69552	2672	55	66814	11.0 M	2476
.org	37621	555	19	37024	6.7 M	2073
.nl	1696103	12304	1002	1682770	233.3 M	1316
.se	334880	8696	100	326067	43.3 M	3681
.uk	10225	314	10	9894	1.6 M	570

Table 15: Overview of DNSSEC data sets

TLD	#domains	#failed	#skipped	#queried	#queries	#auth ns
.com	498502	55909	2231	436593	37.6 M	72168
.net	99564	13904	355	84882	7.4 M	26396
.org	100000	11031	277	88372	7.5 M	27761
.nl	1000000	69092	6812	921441	69.3 M	31108
.se	499999	37361	149560	311871	21.5 M	23756
.uk	26131	3883	92	21858	1.6 M	7091

Table 16: Overview of non-DNSSEC data sets

We ran two scans for each TLD in the source data set. The first scan covered all DNSSEC-signed domains in the TLD. The second scan examined a representative uniformly random sample of unsigned domains with a size in the same order of magnitude as the number of DNSSEC-signed domains in the TLD.

For each scan type Tab. 15 and Tab. 16 show the total number of domains for which queries were attempted, the number of domains for which we failed to obtain the list of authoritative name servers, the number of domains that were skipped (because they were DNSSEC-signed whereas

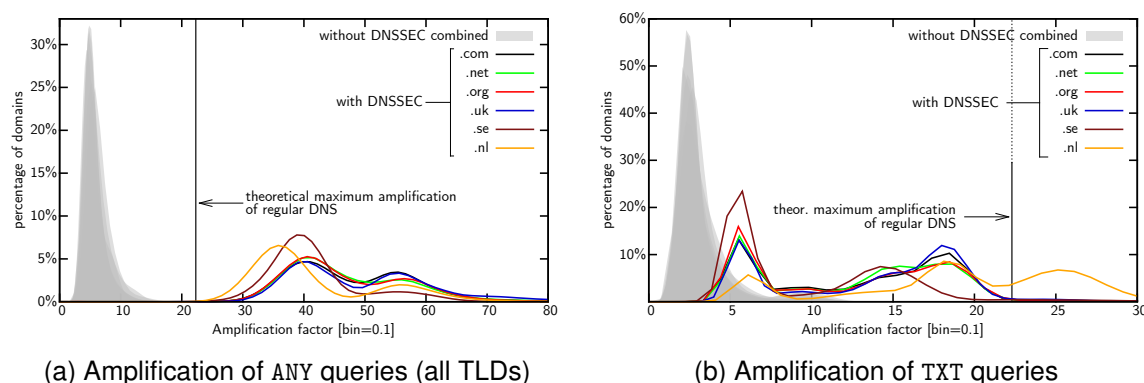


Figure 8: Example for Amplification

they were expected not to be or vice versa), the actual number of domains that were successfully queried, the total number of queries included in the data set and the number of distinct authoritative name servers observed during the scan. We note that there may be a slight difference between the total number of domains for which queries were attempted (col. 2) and the number of failed, skipped and successfully queried domains added up (col. 3 + 4 + 5). This is because for a small number of domains although we were able to determine the set of authoritative name servers none of these responded to queries. Since for both the regular as well as the DNSSEC data sets this difference is very small (0.52% and 0.03% on average over all TLDs respectively) it is not shown in the table.

8.4 Highlights of Results

The comparison of the DNSSEC amplification as measured for the considered query types has depicted a nuanced landscape. As expected, if only ANY queries are considered (see Figure 8a), DNSSEC seems to constitute a serious threat for DDoS reflection and amplification attacks. For this query DNSSEC-signed domains yield high amplification factors, averaging between 40 and 55. This exceeds the average amplification of regular DNS by a factor of $6\times$ - $12\times$. For many common DNS queries, however, using DNSSEC results in larger responses but the amplification factor mostly stays within the acceptable upper limit based on the maximum amplification of classic DNS (see, for example, Figure 8b for TXT queries). Nevertheless, an attacker needs only one or a few domains with large amplification factors, and by carefully choosing a signed domain attackers can achieve significant amplification. It is clear then that this needs to be addressed. The research in [32] points out a number of possible mitigations and countermeasures, among which the large scale deployment of Ingress Filtering (BCP 38 [41]), Response Rate Limiting, EDNS0 cookies and response size limiting.

9 Link Dimensioning

Link dimensioning²² is used by network operators to provision their network links. Typically, the capacity to be provisioned is defined by the average bandwidth utilization added with a safety margin. The bandwidth utilization is commonly obtained by reading interface counters (e.g., via SNMP) every couple of minutes (e.g., 5 to 15 minutes) [43]. The safety margin, on its turn, is determined by rules of thumb such as a percentage of the average bandwidth utilization (e.g., 50% of the calculated bandwidth utilization) [44]. Although this straightforward approach might attend simple requirements of over-provisioning, in more restrictive scenarios it might result in excessive under or over-estimation of actual required capacity for the measured traffic. On the one hand, rough measurement data obtained via counters polling at large timescales (minutes) completely overlook short-term traffic fluctuations. This can result in under-estimation of required capacity during the busy periods, ultimately degrading network performance as perceived by end users. On the other hand, the simplistic rule of thumb to define a safety margin can result in excessive over-estimation of required capacity and, hence, bandwidth resources that could potentially be used for other applications/services are idle and wasted.

The ever increasing demand for bandwidth resources combined with the current trend towards virtualization of network and services will push for more sophisticated link dimensioning approaches to fairly share and allocate bandwidth resources. For instance, the so called Internet big players [45] are already responsible for significant shares of Internet traffic and their shares will soon become even larger. These big players provide essential services to daily life of end users, but most importantly the big players often keep hold of users content within their own datacenters. Retaining users content gives the big players the power to dictate how the Internet should work in the near future. We envision that in the near future, the direct relationship between end users and the big players will narrow. Network operators will still own (most) of the physical infrastructure, but even Internet access will become part of the services hired by users within the big players. This will eliminate the intermediate relationship between end users and network operators. Virtual networks will enable transparent and seamless connections between end users and the big players. However, the coexistence of many virtual networks on top of a single physical infrastructure will bring challenges for proper link resources allocation and usage. In this context, efficient and accurate link dimensioning approaches can certainly make the difference. Such approaches can (i) support operators on the optimal allocation of their link resources, while (ii) ensuring that Quality of Service (QoS) metrics agreed with the big players are met, ultimately (iii) providing end users with good levels of Quality of Experience (QoE).

Focusing on proper allocation of link resources in the future Internet, we have developed and validated a flow-based approach for link dimensioning that is both easy-to-use and accurate. The starting point of our approach is the accurate and already validated link dimensioning formula from [46]. This formula requires traffic statistics that can be easily calculated from packet captures. Although accurate, the drawback of such approach is to require continuous packet capture, which are expensive to do and often demand dedicated hardware/software. Therefore, network operators find it difficult to use due to operational and financial constraints, and they stick to easy-to-use, but inaccurate rules of thumb. Aiming at ease-of-use and accuracy, our approach relies mostly on NetFlow/IPFIX style flow data. Measurement technologies that can provide us with such data are largely found in network devices nowadays and, in many cases, operators already measure flows for other operations such as network monitoring or security. However, flow data imposes challenges for link dimensioning. In particular, the dimensioning formula we use requires a good

²²The research summarized in this section has been published in the paper [42]: R. de O. Schmidt, R. Sadre, A. Sperotto, H. van den Berg, and A. Pras. *A hybrid procedure for efficient link dimensioning*. *Computer Networks*, 67:252-269, 2014.

Table 17: Summary of measurements

abbr.	description	year	length	# of hosts	link capacity	avg. use
A	link from university's building to core router	2011	24h	6.5k	2×1 Gb/s	15%
B	core router of university in The Netherlands	2012	6h	886k	10 Gb/s	10%
C	core router of university in Brazil	2012	84h45min	10.5k	155 and 40 Mb/s	19%
D	backbone links connecting Chicago and Seattle	2011	4h	1.8M	2×10 Gb/s	8%
E	backbone links connecting San Jose and Los Angeles	2011–2012	5h	3M	2×10 Gb/s	10%
F	trans-Pacific backbone link	2012	13h15min	4M	n/a	n/a

estimation of traffic variance. Given that in flows we miss information of individual packets, i.e., we do not know their respective sizes and how they are distributed in time, it is difficult to estimate traffic variance. Our proposed approach relies on flow data combined with mathematical models to estimate traffic variance and, ultimately, to estimate the required capacity at timescales as low as 1ms. This approach has been validated on a set of traces collected worldwide, which we summarize in Table 17. The entire dataset comprises 548 15-minute traces totaling 137 hours of captures. The trace duration of 15 minutes has been chosen in accordance with [47] such as to comply with a stationarity assumption. These traces come from different locations around the globe and account for a total of more than 13.3 billion packets. Traffic captures were done at the IP packet level, using tools such as `tcpdump`.

In Figure 9, we use the empirical estimation of required capacity C_{emp} to validate the estimation of required capacity C_{flow} from our flow-based approach. This empirical capacity is manually determined from the packet traces we used to validate our flow-based approach for link dimensioning.

Figure 10 gives an idea of how the estimation compares to the actual traffic stream. This figure shows the traffic throughput of a 15-minute trace at timescale of 10ms (time bins in the x-axis). The straight lines show the estimations C_{flow} and C_{emp} , and the one obtained using rule of thumb C_{RoT} (50% of average utilization). Note that in few time bins the traffic throughput actually exceeds the estimated capacity C_{flow} . That is because the dimensioning formula we use allows for the definition of an exceedance probability ε . This value tells the formula the fraction of bins in which traffic throughput can be higher than C_{flow} , assuming that buffers are capable to handle exceeding traffic. For the examples of Figures 9 and 10, ε was set to 0.01 (i.e., accepting for 1% of time bins to be underestimated). Clearly, the estimation of required capacity C_{flow} is successful when the fraction of exceeding time bins $\hat{\varepsilon}$ is less than or equal to the defined as the acceptable fraction ε (i.e., $\hat{\varepsilon} \leq \varepsilon$). Notice that, this is equivalent of the required capacity C_{flow} being greater than or equal to the minimum capacity determined by the empirical C_{emp} (i.e., $C_{flow} \geq C_{emp}$). It is important to understand in Figure 10 that C_{flow} from our approach underestimates the required capacity in this sample trace, because the obtained exceedance is $\hat{\varepsilon} = 0.0127$. Nonetheless, one can clearly see that the simple rule of thumb resulted in a C_{RoT} that excessively overestimates the actual required capacity expressed by C_{emp} , which might ultimately lead to the allocation of unnecessary link resources.

The proposed flow-based approach was extensively validated using hundreds of traffic traces, comprising many hours of captures, from six different locations around the globe, and captured

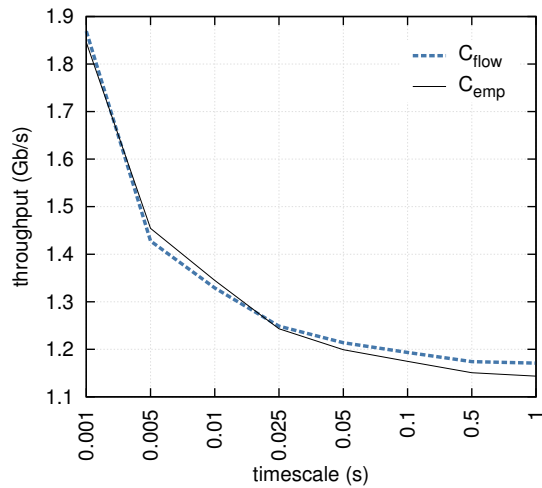


Figure 9: Estimation of required capacity at various timescales for a sample trace; exceedance probability $\varepsilon = 0.01$.

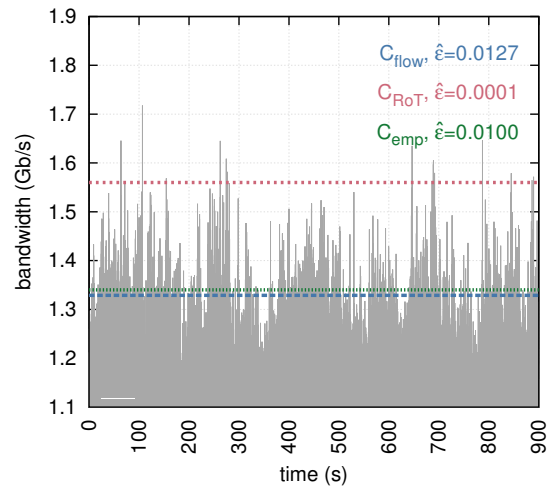


Figure 10: Estimations of required capacity as compared to the actual traffic throughput for a sample trace; timescale of 10ms and exceedance probability $\varepsilon = 0.01$.

at universities and ISP backbone links. The complete validation and set of results can be found in [42].

10 Integration of EU research

The overall FLAMINGO activities in the context of the integration of the EU research landscape are reported in D3.3. 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 have 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 **Ethics in Data Sharing**²³ (Schloss Dagstuhl, January 26-30, 2014). The seminar was attended by a variegated group of researchers from the disciplines of computer science, ethics, and law, brought together by background experiences in data sharing. The major outcome of the seminar has been a collaboration between UT, SURFnet bv, University of Amsterdam, Tilburg University and UZH (also explored in WP7) that aims at developing a usable policy for facilitating ethical data sharing between data producer (e.g., operators) and data consumers (e.g., researchers). This activity is strongly backed up by SURFnet, the Dutch NREN, which will run a first pilot of the proposed policy. The seminar was organized by S. Dietrich (Stevens Institute of Technology, US), M. Hildebrandt (Free University of Brussels, BE), Aiko Pras (UT) and Lenore D. Zuck (University of Chicago, US).
- The Dagstuhl seminar **Global Measurement Framework**²⁴ (Schloss Dagstuhl, November 17-20, 2013). The seminar brought together researchers from industry, academia, and regulators across continents and across different backgrounds to discuss the state of the art in measurements and their exploitation, measurement and analysis techniques, privacy and anonymization. The seminar was organized, in collaboration with the EU projects Leone²⁵ and mPlane,²⁶ by: Philip Eardley (BT Research, GB), Marco Mellia (Polytechnic University of Torino, IT), Jörg Ott (Aalto University, FI), Jürgen Schönwälder (Jacobs University Bremen, DE) and Henning Schulzrinne (Columbia University, US).
- Anna Sperotto (UT), Marinos Charalambides (UCL) and Jeroen Famaey (iMinds) have taken part in the organization of the **8th International Conference on Autonomous Infrastructure, Management and Security** (AIMS 2014, Brno, Czech Republic), as Conference Co-chair, PhD Student Workshop Co-chair and Lab Co-Chair, respectively. In addition, several members of the consortium have acted as TPC members.
- The **6th Workshop on the Usage of NetFlow/IPFIX in Network Management** is currently planned to be co-located with a Network Management Research Group (NMRG) meeting at IETF 92 in Prague (July 2015). 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.

²³<http://www.dagstuhl.de/14052>

²⁴<http://www.dagstuhl.de/13472>

²⁵ <http://www.leone-project.eu/>

²⁶<http://www.ict-mplane.eu/>

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 Leone²⁵ – From global measurements to local management (grant no. 317647).
- The FP7 Project mPlane²⁶ – Building an Intelligent Measurement Plane for the Internet (grant no. 318627).
- The FP7 Project SmartenIT²⁷ – Socially-aware Management of New Overlay Application Traffic combined with Energy Efficiency in the Internet (grant no. 317846).
- The FP7 Large-scale Integrating project Mobile Cloud Networking²⁸ (grant no. 318109)
- The FP7 project EVANS²⁹ (grant no. PIRSES-GA-2010-269323)
- the FP7 project ALIEN³⁰ – Abstraction Layer for Implementation of Extensions in Programmable Networks (grant no. 317880).

FLAMINGO, Leone and mPlane contributed to the organization of the Dagstuhl seminar **Global Measurement Framework** (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.

²⁷<http://www.smartenit.eu/>

²⁸<http://www.mobile-cloud-networking.eu>

²⁹<http://www.fp7-evans.eu/>

³⁰<http://www.fp7-alien.eu>

11 Conclusions

This deliverable described WP5 achievements in the field of network and service monitoring. 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 show active research in the WP-specific objectives, among which research topics focusing on security are particularly active.

The researchers – answering a comment of the reviewers – have also targeted high-impact conferences and journals in the larger areas of networking and network monitoring, which resulted in publications at INFOCOM 2014, ACM IMC 2014, ACM Computer Communication Review, TNSM, IEEE Surveys & Tutorials and Elsevier Computer Networks. In addition, several research activities lead to the publication of datasets, which are a strong contribution to the community.

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 50 papers, of which several have a strong measurement component.

Finally, also in Y2 collaborations are key in the project approach to research. As such, the PhD collaborations remain one of the major achievement of the project. These collaborations are also the basis of the strong collaboration among the research work packages. In Y2 WP5 was involved in a lively group of PhD collaborations. Several PhD collaborations have continued from Y1, while other have started in Y2. Compared to Y1, this year the PhD collaborations have also had a strong scientific output, and 9 publications are authored by more than one FLAMINGO partner. In addition to the Ph.D. collaborations, WP5 has also been involved in collaborations with other EU projects and institutions. We believe that good collaborations are one of the components of the outstanding scientific output of the workpackage, and WP5 plans to continue supporting and encouraging collaborations also in Y3.

Abbreviations

AI	Artificial Intelligence
API	Application Programming Interface
BSD	Berkeley Software Distribution
CDN	Content Delivery Network
CPS	Cyber Physical Systems
DDoS	Distributed Denial of Service attack
DNS	Domain Name System
DNSSEC	Domain Name System Security Extensions
DODAG	Destination Oriented Directed Acyclic Graph
EDNS	Extension mechanisms for DNS
FN	False Negatives
ForCES	Forwarding and Control Element Separation
FP	False Positives
Gbps	Giga bits per second
HAS	HTTP Adaptive Streaming
HTTP	Hyper-text Transfer Protocol
HTTPS	Hyper-text Transfer Protocol Secure
ICMP	Internet Control Message Protocol
ID	Intrusion Detection
IDS	Intrusion Detection System
ILP	Integer Linear Program
IP	Internet Protocol
IPFIX	Internet Protocol Flow Information Export
ISP	Internet Service Provider
LLN	Low-power and Lossy Networks
LRU	Least Recently Used
MNO	Mobile Network Operator
MTR	Multi-Topology Routing
NFQL	Network Flow Query Language

NMRG	Network Management Research Group
NREN	National Research and Education Network
OS	Operative System
PPF	Packets-per-flow
QoS	Quality-of-Service
RPL	Routing Protocol for Low power and Lossy Networks
RTT	Round-Trip Time
SDN	Software-Defined Networking
S.M.A.R.T	Specific Measurable Achievable Relevant Timely
SNMP	Simple Network Management Protocol
SSH	Secure SHell
TCP	Transmission Control Protocol
TLD	Top-Level Domain
TN	True Negatives
TP	True Positives
TPM	Trusted Platform Module
UDP	User Datagram Protocol
VLAN	Virtual Local Area Network
VP	Virtual Network
VNP	Virtual Network Provider
VoIP	Voice-over-IP
VoD	Video on Demand
VoS	Value of Service

A Internet Traffic Statistics

This appendix includes the poster, titled “ReFlow – Statistics on Internet Traffic” by M. Hoogesteger, R. Schmidt, A. Sperotto and A. Pras (UT), that has been presented at the Terena Networking Conference 2014 [48].

The goal of the poster is twofold. On the one side it reports to the TNC community about the progress of the *Internet Traffic Statistics* project, which has already been presented at the TERENA Networking Conference (TNC) 2013. On the other side, it aimed at maintaining and raising awareness about the project and attracting other data providers, such as other Internet Service Providers.

The *Internet Traffic Statistics* presents an extended version of the Internet's traffic weekly reports and aims at collecting traffic statistics from operators around the world, and make them available in a public repository with a friendly user interface. Traffic data will be collected at, for example, NREs and universities, allowing us to have a broad picture of the Internet usage. Traffic statistics will be generated from flow data (e.g., NetFlow). The poster provides an overview of the current data inflow to the process as well as it provides a first example of trend analysis that can be conducted with the data stored in the system.

ReFlow

Statistics on Internet Traffic

Goal

- Create an online repository containing statistics about Internet traffic, which will serve as a reference to network operators and the research community.

ReFlow operation

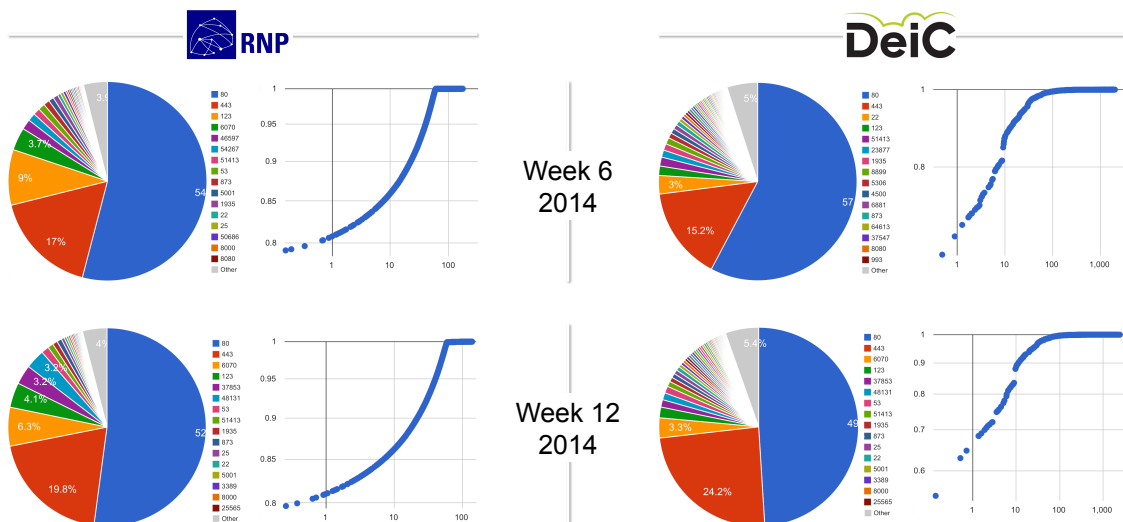
- The tool processes NetFlow-like traffic measurements and computes traffic statistics that later are presented in a public website.



stats.simpleweb.org

Present

- Currently, we are populating the repository with statistics of traffic from locations in Europe and South America. We have a fully operational system, running at our side and remote scripts at collaborators.




Future

- We aim at gathering data from other sources to become as representative as possible of general Internet traffic.
- Create an interface for more flexible querying of statistics, allowing for the comparison of data between various sources and weeks.

Do you have data? Contact us!

- Do you have NetFlow measurements or similar?
- Data can be processed at either our side or yours.
- We **do not** use or disclose sensitive data.



Martijn Hoogesteger, Ricardo de O. Schmidt, Anna Sperotto and Aiko Pras
Design and Analysis of Communication Systems
m.hoogesteger@student.utwente.nl, r.schmidt@utwente.nl

UNIVERSITY OF TWENTE.

SURF NET



B Towards Comparability of Intrusion Detection Systems: New Data Sets

This appendix includes the poster, titled “Towards Comparability of Intrusion Detection Systems: New Data Sets” by R. Koch, M. Golling and G. Dreo Rodosek (UniBwM), that has been presented at the TERENA Networking Conference (TNC) 2014 [49].

Contemporaneously with the overwhelming success of the Internet and its dissemination in all areas of life, attacks have risen dramatically over the past few years - both from a qualitative as well as a quantitative point of view. As a consequence, despite other mechanisms, intrusion detection systems (IDSs) are under intense investigation. Different systems have been put on the market and numerous research contributions continuously try to improve the performance of the systems w.r.t. detection and false alarm rates as well as data rates that can be processed. Despite more than 30 years of research, the practical deployment of IDSs still suffers from substantial weaknesses: Knowledge-based systems are no longer in a good position to withstand the latest malware adequately. The scope of the signatures used is increasing more and more throughout the past years.

However, this in turn solves the real problem less and less. Firstly, a signature must initially be developed for the malware, delaying the detection capability and secondly - as a study by McAfee shows impressively - these signatures are less successful. On average, signatures are taking effect less than 10 times worldwide. On the other hand, behaviour-based systems are still affected by high false alarm rates, which complicates a practical usage. Even more, an adequate evaluation of the proposed improvements for IDSs also represents a major challenge. Even today, evaluations are often based on the DARPA 98/99 data sets. However, these data sets have several shortcomings and thus have been criticized a lot. Therefore, researches advised to stop using them any longer for evaluation purposes. In the meantime, several other data sets have been released, for example a redesign of the DARPA data set, the data of the MAWI Working Group, the MoMe Cluster, data from the Consortium Internet 2, from ACM SIGCOMM, from the Cooperative Association for Internet Data Analysis (CAIDA), RIPE or the Internet Archive. Unfortunately, none of them was able to get accepted comprehensively throughout the community: Specific scenarios, limited availability, a lack of ground truth, etc., prevent a wide application and acceptance.

To solve this issue, the corresponding poster aims to encourage Internet Service Providers (ISPs) and/or companies (SMB as well as large companies) to create a new, realistic data set in order to set the basis for a resilient evaluation of IDSs. In contrast to works of others, we aim to produce a data set that is realistic, up to date and generally applicable. With regard to backbone operators, we want to receive real flow data, while for company networks, we propose to collect flow data as well as full captures. For legal and privacy reasons, all data will be pseudonymized. In addition, we plan to keep this dataset updated by producing a new release twice per year.

Towards Comparability of Intrusion Detection Systems: New Data Sets

Robert Koch, Mario Golling and Gabi Dreo Rodosek

Universität der Bundeswehr München
 {robert.koch, mario.golling, gabi.dreo}@unibw.de

Problem Statement

- ▶ Attacks have risen dramatically over the past few years - both from a qualitative as well as a quantitative point of view
- ▶ Throughout the years, different Intrusion Detection Systems (IDSs) have been put on the market and numerous research contributions continuously try to improve the performance of the systems wrt. detection and false alarm rates as well as data rates
- ▶ However, even up to now, an **adequate evaluation** of the proposed improvements for **IDSs represents a major challenge**
- ▶ Evaluations are often based on the **DARPA 98/99 data sets**
- ▶ However, these data sets have several shortcomings and thus have been criticized a lot; e.g.:
 - ▶ Mahoney and Chan discovered that all malicious data packets of the DARPA set have a time to live (TTL) of 126 or 253, whereas the benign packets of the background traffic mainly have a TTL of 127 or 254
 - ▶ The malicious TTL values 126 and 253 do not occur in the (*attack-free*) training records
 - ▶ In addition, the TTL values of the records are, on average, significantly higher compared to real-world traffic
- ▶ Therefore, **researches advised to stop using the DARPA 98/99 data sets any longer for evaluation purposes**
- ▶ Other data sets were not able to become accepted

Requirements

- ▶ **General Statistical Requirements**
 - ▶ Inferring from the sample (= data set) to the overall population can only be meaningful, when the **sample reflects all characteristics of the population as good as possible** (i.e., if the sample is representative and if there is a structural equality between sample and population)
 - ▶ Within the **stratified random sampling method** the corresponding **population is divided into several smaller groups**, which are called layers
 - ▶ Then, samples are chosen separately
 - ▶ It is also important that the information (at the different layers) are **up to date**
- ▶ **Specific Technical Requirements**
 - ▶ As Network Intrusion Detection Systems (NIDSs) operate according to different principles (in particular **payload-, behavior- or flow-based**), a data set must be present for an evaluation of both
 - ▶ Since a synthetic record often constrains the general statistical requirements, a data set should therefore be created of **real-world data** (although this usually means that the data - here especially the IPs - must be made anonymous / pseudonymous)
 - ▶ A data set should cover all relevant aspects and therefore should contain **2-3 typical "operating cycles"** (in the case of company networks) or **at least one month**


Characteristics of Available Data Sets

DATA SET	PAYLOAD	FLOWS	SYNTHETIC GENERATION	REAL DATA	CURRENT ATTACKS	GROUND TRUTH	DATASET AVAILABILITY	ONGOING/ UPDATED
DARPA 98/99		X	✓	X	X	✓	✓	X
KDD 99	✓	X	✓	X	X	✓	✓	X
DARPA REDESIGN	✓	X	✓	X	(✓)	✓	X	X
MAWI WORKING GROUP	✓	X	X	✓	(✓)	X	✓	✓
MOME CLUSTER	✓	(✓)	X	✓	(✓)	X	✓	(X)
CONSORTIUM INTERNET 2	X	X	X	✓	(✓)	X	✓	✓
LABELLED FLOWS	X	✓	X	(✓)	✓	(✓)	✓	X
CAIDA	✓	X	X	✓	(✓)	X	(✓)	✓
WAIKATO	(✓)	X	X	✓	(✓)	X	(✓)	X
RIPE	(✓)	X	X	✓	X	X	(✓)	(X)
INTERNET ARCHIVE	X	X	X	✓	X	X	✓	X
UMASS	✓	X	(✓)	(✓)	(✓)	X	✓	X
PREDICT	X	✓	✓	✓	(✓)	X	(✓)	(✓)

- ▶ Other lesser-known data sets, e.g., from the University of Aveiro or Crawdad, are not listed because of not being useful for IDSs or not being available
- ▶ (...) signifies limitations, e.g., only some attack classes are detectable, registration is required or only parts of the data is updated

Proposal for Cooperation

- ▶ Due to the shortcomings of available data sets, **this poster aims to encourage** Internet Service Providers (ISPs) and/or companies (SMB as well as large companies) **to collaborate on the creation of a new, realistic data set in order to lay the foundation for a resilient evaluation of IDSs**
- ▶ In contrast to works of others, **we aim to produce a data set that is realistic, up to date and generally applicable**
- ▶ Following the idea of the **stratified random sampling method**, data needs to be obtained from different layers of the Internet (ISP/Backbone Operator, Companies, Academic Networks, End Customer, etc.)



- ▶ With regard to **backbone operators**, we like to receive real **flow data**, while for **company networks**, we propose to collect **flow data as well as full captures**
- ▶ For legal and privacy reasons, all data will be **pseudonymized**
- ▶ In addition, we plan to keep this data set updated by producing **a new release twice a year**

⇒ **Contact us, if you can participate in a joint continuous generation of a real-world IDS data set** ⇐

C SSHCure: SSH Intrusion Detection using NetFlow and IPFIX

This appendix includes the poster, titled “SSHCure: SSH Intrusion Detection using NetFlow and IPFIX” by L.Hendriks, R.Hofstede, A. Sperotto and A. Pras (UT), that has been presented at the TERENA Networking Conference (TNC) 2014 [50].

SSHCure is the first IDS capable of distinguishing successful attacks from unsuccessful attacks, thus detecting actual compromises. As powerful as SSH is to administrators, as attractive it is to anyone with malicious intents. Measurements showing more than 700 attacks on NRENs per day emphasize this. This number is also the source of the main problem in existing detection systems: while 699 of these attacks are typically unsuccessful and therefore not interesting to network administrators or CSIRT members, a single successful one is. And its consequences possibly include severe damage to the target hosts themselves, others hosts in the network, or even the network itself: an NREN should be informed as quickly as possible when this happens, so adequate actions can be undertaken.

In SSHCure, R. Hofstede and L.Hendriks implement a detection algorithm based on flow export technologies, i.e. NetFlow and IPFIX. A flow-based approach brings clear performance benefits over packet-based approaches in large-scale networks. The packet payloads are not available in flow data, making it more privacy preserving, while the loss of information (in comparison to a packet-based approach) is limited due to the encrypted nature of SSH. We show however, that flow data bring sufficient information to perform accurate detection. Moreover, flow export technologies are widely available on high-end networking devices. SSHCure is a plugin for NfSen (a flow collector for NetFlow and IPFIX, used by many in the NREN community) and therefore easy to install and use within all kinds of networks. The adoption of SSHCure underlines this, as it is currently deployed at several large commercial ISPs, CERTs and NRENs. All of these types of organizations need to be able to act swiftly when a compromise has been observed, and SSHCure is designed to support in that: the web-interface gives clear insight on the situation, including detailed information on both attacker and targets, comprehensible visualisations of network flows, and raw flow data for extensive analysis if needed. This is backed up by a flexible notification system, and (currently under development) integration with incident reporting systems via standard protocols (e.g. IODEF or X-ARF).


SSHCure, available via Sourceforge [51] and recently moved to GitHub³¹ has been in development for 2.5 years, and is still actively being developed and supported. The first prototype was presented at the Autonomous Infrastructure, Management and Security conference (AIMS) in 2012 [15], and promising results were achieved. With the latest available version, we performed extensive validation using datasets from both campus and backbone networks. Results show detection rates of up to 100%.

³¹<https://github.com/SSHCure/SSHCure>)

SSHCure: SSH Intrusion Detection using NetFlow and IPFIX

Luuk Hendriks, Rick Hofstede, Anna Sperotto, Aiko Pras

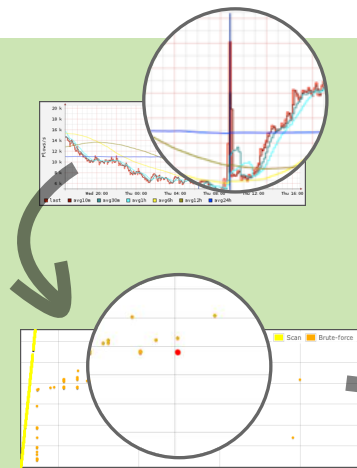
Design and Analysis of Communication Systems, University of Twente
{luuk.hendriks,r.j.hofstede,a.sperotto,a.pras}@utwente.nl



We see an **increasing number of SSH attacks**. Measurements show that a typical NREN enables for **more than 1000 attacks per day**. While few of these are successful (i.e. end in a **compromise**), the ones that do succeed can cause **severe damage** in a plethora of different ways.

- Scalable, deployed in NREN with up to 13 backbone links
- Privacy preserving
- No performance impact on the network level
- Accurate detection for different phases within an attack
- Easily deployed as an NfSen plugin

SSHCURE



SSHCure is able to analyze large amount of flow data and show what is really going on in the network, alerting administrators in real time.



Adoption and deployments:

 Hogeschool van Amsterdam
Media, Creatie en Informatie

 CESNET

 SURF NET

UNIVERSITY OF TWENTE.

Get SSHCure at <http://sshcure.sf.net>

 SURF NET

UNIVERSITY OF TWENTE.




D Real-time DDoS Defense: A collaborative Approach at Internet Scale

This appendix includes the poster, titled “Real-time DDoS Defense: A collaborative Approach at Internet Scale” by Jessica Steinberger (University of Applied Sciences Darmstadt and UT), Anna Sperotto (UT), Aiko Pras (UT) and Harald Baier (University of Applied Sciences Darmstadt), that has been presented at the TERENA Networking Conference (TNC) 2014 [52]. The poster has won the **TNC 2014 best student poster** competition.

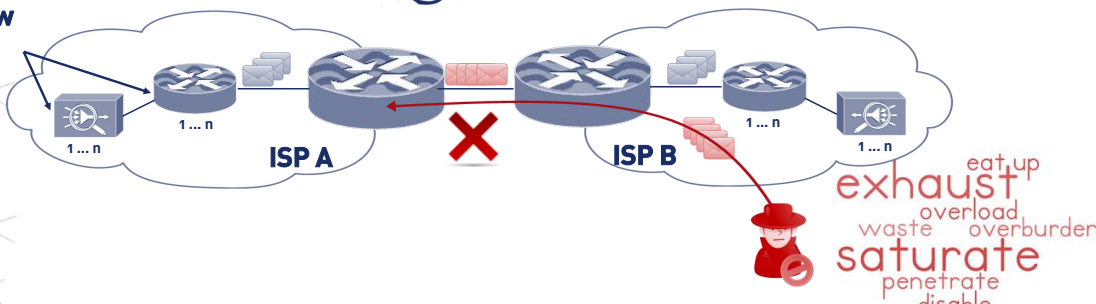
In the last years, Distributed Denial of Service attacks (DDoS) evolved to one of the major causes responsible for network infrastructure and service outages. Often the amount of traffic generated by DDoS attacks is such that, although traditional security solutions as firewalls and Intrusion Prevention Systems are deployed, the target network will lose connectivity, because the network resources are exhausted. To optimize mitigation and response capabilities and thus reduce potential damages caused by DDoS attacks, mitigation and response should be moved from the target network to the networks of Internet Service Providers (ISPs). Additionally, ISPs should collaborate and exchange information in context of network security. The poster proposes a framework for flow-based real-time and automatic mitigation of DDoS attacks in ISP networks. The framework collects and processes network flow-based data e.g. NetFlow/IPFIX from the network edge router of an ISP network. The collected data is used to perform anomaly detection, data fusion and classification. In case of a detected anomaly within the flow-based data a security event is raised. Based on this security event a collaborative process is initiated. The framework collaborates with third parties by gathering and processing security information e.g. from other ISPs, customers or available data e.g. Blacklists, Open DNS resolvers etc.).

Real-time DDoS Defense

A Collaborative Approach

Problem:
 What happens, if 400 Gbps are reaching  network? [1][2]

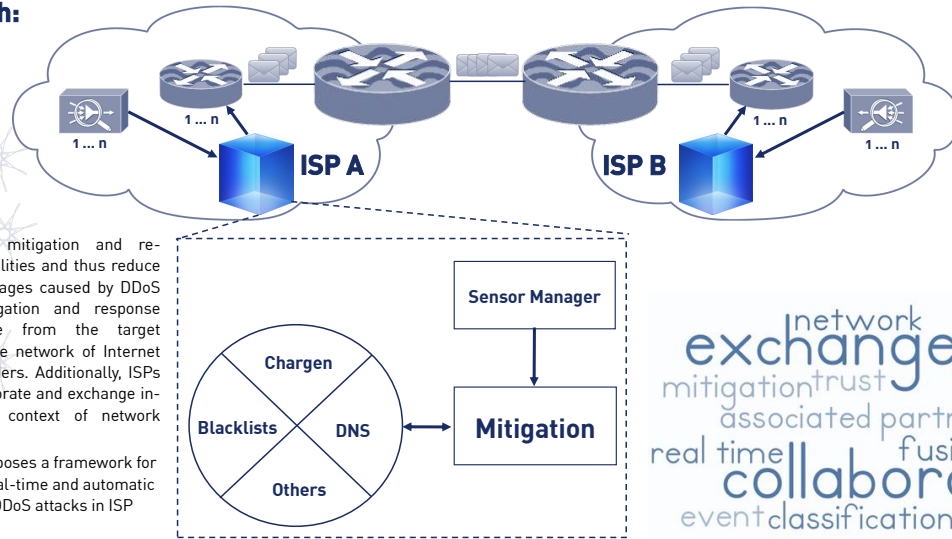
NetFlow
IPFIX



Research Questions:

1. Is real-time and automatic mitigation at ISP level performed and if yes, how?
2. How can the effect of DDoS attack be limited?
3. How can the framework for real-time and automatic mitigation be validated?

Approach:



To optimize mitigation and response capabilities and thus reduce potential damages caused by DDoS attacks, mitigation and response should move from the target network to the network of Internet Service Providers. Additionally, ISPs should collaborate and exchange information in context of network security.

This work proposes a framework for flow-based real-time and automatic mitigation of DDoS attacks in ISP networks.

network
 exchange
 mitigation trust response
 associated partner
 real time collaboration fusion
 event classification

[1] Anstee, D., Bussiere, D., Sockrider, G., Morales, C.: Worldwide Infrastructure Security Report. Technical Report IX, Arbor Networks Inc. (January 2013) <http://www.arbornetworks.com/research/infrastructure-security-report>.

[2] Prince, M. Technical Details behind a 400 Gbps NTP Amplification DDoS attack (February 2014) <http://blog.cloudflare.com/technical-details-behind-a-400gbps-ntp-amplification-ddos-attack>

References

- [1] G. van den Broek, R. van Rijswijk-Deij, A. Sperotto, and A. Pras. DNSSEC meets real world: dealing with unreachability caused by fragmentation. *IEEE communications magazine*, 52(4):154–160, April 2014.
- [2] Rick Hofstede, Luuk Hendriks, Anna Sperotto, and Aiko Pras. SSH Compromise Detection using NetFlow/IPFIX. *ACM SIGCOMM Computer Communication Review*, 44(5):21–26, 2014. (to appear).
- [3] Maxim Claeys, Steven Latré, Jeroen Famaey, and Filip De Turck. Design and Evaluation of a Self-Learning HTTP Adaptive Video Streaming Client. *Communications Letters, IEEE*, 18(4):716–719, April 2014.
- [4] G. Hurel, R. Badonnel, A. Lahmadi, and O. Festor. Outsourcing Mobile Security in the Cloud. In *Proc. of the 8th International Conference on Autonomous Infrastructure, Management and Security (AIMS 2014)*, pages 69–73. 2014.
- [5] I. Drago, R. de Oliveira Schmidt, R. J. Hofstede, A. Sperotto, M. Karimzadeh Motallebi Azar, B. R. H. M. Haverkort, and A. Pras. Networking for the Cloud: Challenges and Trends. *PIK - Praxis der Informationsverarbeitung und Kommunikation*, 36(4):207–214, December 2013.
- [6] F. Idzikowski, S. Orlowski, C. Raack, H. Woesner, and A. Wolisz. Saving energy in IP-over-WDM networks by switching off line cards in low-demand scenarios. In *Proceedings of the 14th conference on Optical Network Design and Modeling (ONDM'10)*, pages 42–47, 2010.
- [7] A. Coiro, F. Iervini, and M. Listanti. Distributed and adaptive interface switch off for internet energy saving. In *Proceedings of the 20th International Conference on Computer Communications and Networks (ICCCN)*, pages 1–8, 2011.
- [8] M. Charalambides, D. Tuncer, L. Mamatas, and G. Pavlou. Energy-aware adaptive network resource management. In *Integrated Network Management (IM 2013), 2013 IFIP/IEEE International Symposium on*, pages 369–377, May 2013.
- [9] D. Tuncer, M. Charalambides, G. Pavlou, and N. Wang. Dacorm: A coordinated, decentralized and adaptive network resource management scheme. In *Network Operations and Management Symposium (NOMS), 2012 IEEE*, pages 417–425, April 2012.
- [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. 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] G. Dreo, M. Golling, W. Hommel, and F. Tietze. ICEMAN: An architecture for secure federated inter-cloud identity management. In *IFIP/IEEE International Symposium on Integrated Network Management (IM 2013)*, May 2013.
- [18] G. Dreo, M. Golling, and W. Hommel. MuSIC: An IT security architecture for inter-community clouds. In *IFIP/IEEE International Symposium on Integrated Network Management (IM 2013)*, May 2013.
- [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] R. Mijumbi, J. Serrat, J.L. Gorricho, M. Claeys, F. De Turck, and S. Latré. Design and Evaluation of Learning Algorithms for Dynamic Resource Management in Virtual Networks. In *14th IEEE/IFIP Network Operations and Management Symposium (NOMS 2014) (submitted to)*.
- [25] Rashid Mijumbi, Juan-Luis Gorricho, and Joan Serrat. Learning algorithms for dynamic resource allocation in virtualised networks. In *Management of Large Scale Virtualized Infrastructures: Smart Data Acquisition, Analysis and Network and Service Management in the Future Internet*, 2014.
- [26] Rashid Mijumbi, Juan-Luis Gorricho, and Joan Serrat. Contributions to efficient resource management in virtual networks. In *IFIP 8th International Conference on Autonomous Infrastructure, Management and Security (AIMS), 2014*, 2014.
- [27] Rashid Mijumbi, JL Gorricho, J Serrat, Maxim Claeys, Jeroen Famaey, and Filip De Turck. Artificial neural network-based autonomous allocation of resources in virtual networks. In *To appear in proceedings of the European Conference on Networks and Communications (EUCNC)*, 2014.
- [28] R. Mijumbi, J. Serrat, J. L. Gorricho, M. Shen, K. Xu, and K. Yang. A neuro-fuzzy approach to self-management of virtual network resources. In *Journal of Expert Systems With Applications (submitted to)*, 2014.
- [29] Rashid Mijumbi, J Serrat, J Rubio-Loyola, Niels Bouten, Filip De Turck, and Steven Latré. Dynamic resource management in sdn-based virtualized networks. In *1st International Workshop on Management of SDN and NFV Systems*, 2014.

- [30] 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-01>, July 2014.
- [31] 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.
- [32] R. van Rijswijk-Deij, A. Sperotto, and A. Pras. DNSSEC and its potential for DDoS attacks. In *Proceedings of the ACM Internet Measurement Conference 2014 (IMC 2014)*, 2014.
- [33] J. Araujo, R. Landa, R. Clegg, G. Pavlou, and K. Fukuda. A longitudinal analysis of Internet rate limitations. In *Proceedings of IEEE International Conference on Computer Communications (INFOCOM 2014)*, 2014.
- [34] K. Cho, K. Mitsuya, and A. Kato. Traffic data repository at the WIDE project. In *Proceedings of the USENIX Annual Technical Conference*, 2000.
- [35] Venafi, Inc. Ponemon 2014 SSH Security Vulnerability Report. Technical report, Venafi, 2014.
- [36] Anna Sperotto, Ramin Sadre, Pieter-Tjerk de Boer, and Aiko Pras. Hidden Markov Model modeling of SSH brute-force attacks. In *Proceedings of DSOM'09*, 2009.
- [37] Laurens Hellemons, Luuk Hendriks, Rick Hofstede, Anna Sperotto, Ramin Sadre, and Aiko Pras. SSHCure: A Flow-Based SSH Intrusion Detection System. In *Proceedings of AIMS'12*, 2012.
- [38] P. Vixie. Extension Mechanisms for DNS (EDNS0). RFC 2671 (Proposed Standard), August 1999. Obsoleted by RFC 6891.
- [39] P. Mockapetris. RFC 1035 - Domain Names - Implementation and Specification, 1987.
- [40] M. Kühner, T. Hupperich, C. Rossow, and T. Holz. Exit from Hell? Reducing the Impact of Amplification DDoS Attacks. In *Proc. of the 23rd USENIX Security Symposium*, August 2014.
- [41] P. Ferguson and D. Senie. BCP 38 - Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing, 2000.
- [42] R. de O. Schmidt, R. Sadre, A. Sperotto, H. van den Berg, and A. Pras. A hybrid procedure for efficient link dimensioning. *Computer Networks*, 67:252 – 269, 2014.
- [43] Cisco Systems Inc. How To Calculate Bandwidth Utilization Using SNMP. http://www.cisco.com/image/gif/paws/8141/calculate_bandwidth_snmp.pdf, 2005. Online. Accessed Apr. 2014.
- [44] Cisco Systems Inc. Best Practices in Core Network Capacity Planning. http://www.cisco.com/c/en/us/solutions/collateral/service-provider/quantum/white_paper_c11-728551.pdf, 2013. Online. Accessed Aug. 2014.
- [45] Vinicius Gehlen, Alessandro Finamore, Marco Mellia, and Maurizio M. Munafò. Uncovering the Big Players of the Web. In *Proceedings of the 4th International Workshop on Traffic Monitoring and Analysis, TMA'12*, pages 15–28, 2012.
- [46] Remco van de Meent. *Network Link Dimensioning: A Measurement & Modeling Based Approach*. PhD thesis, University of Twente, 2006. ISSN 1381-3617.

- [47] R. van de Meent, M. Mandjes, and A. Pras. Gaussian traffic everywhere? In *Communications, 2006. ICC '06. IEEE International Conference on*, volume 2, pages 573–578, June 2006.
- [48] M. Hoogesteger, R. de O. Schmidt, A. Sperotto, and A. Pras. ReFlow - Statistics on Internet Traffic (Poster). TERENA Networking Conference (TNC 2014).
- [49] R. Koch, M. Golling, and G. Dreo Rodosek. Towards Comparability of Intrusion Detection Systems: New Data Sets (Poster). TERENA Networking Conference (TNC 2014).
- [50] L. Hendriks, R. Hofstede, A. Sperotto, and A. Pras. SSHCure: SSH Intrusion Detection using NetFlow and IPFIX (Poster). TERENA Networking Conference (TNC 2014).
- [51] R. Hofstede and L. Hendriks. SSHCure – Flow-based SSH Intrusion Detection System. <http://sourceforge.net/projects/sshcure/>, Sept. 2013.
- [52] J. Steinberger, A. Sperotto, A. Pras, and H. Baier. Real-time DDoS Defense: A collaborative Approach at Internet Scale (Poster). TERENA Networking Conference (TNC 2014).