Industry-Academia Open Research Collaboration: A Small Tech Company Perspective (SIDN)

Giovane C. M. Moura
SIDN Labs and TU Delft
The Netherlands
giovane.moura@sidn.nl

Ralph Koning
SIDN Labs and University of Amsterdam
The Netherlands
ralph.koning@sidn.nl

Moritz Müller
SIDN Labs and University of Twente
The Netherlands
moritz.muller@sidn.nl

Cristian Hesselman
SIDN Labs and University of Twente
The Netherlands
cristian.hesselman@sidn.nl

This article is an editorial note submitted to CCR. It has NOT been peer reviewed.
The authors take full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

Multiple studies have shown the benefits of industry-academia collaborations. In this editorial, we discuss the industry-academic research collaboration from the perspective of SIDN, a Netherlands-based privately held small tech company with a public-interest mission, which operates the `.nl` top-level domain. We present and compare five models for industry-academia collaboration and show how each has produced deliverables that benefit not only SIDN and its partners involved but also the broader community, including the SIGCOMM's. We also share lessons learned from building and running an in-house research team. Our goal is to encourage other industry players to pursue open, collaborative research that serves academia, industry, and society alike.

KEYWORDS

Industry, Academia, Collaboration, DNS Registry, Open Science

ACM Reference Format:
Giovane C. M. Moura, Ralph Koning, Moritz Müller, and Cristian Hesselman. 2025. Industry-Academia Open Research Collaboration: A Small Tech Company Perspective (SIDN). In *ACM SIGCOMM Computer Communication Review, Volume 55, Issue 3, July 2025*. ACM, New York, NY, USA, 6 pages.

1 INTRODUCTION

A previous CCR editorial presented research models for academia to collaborate with Google [41]. It showed how academics can access Google's networking datasets given privacy, costs, and business constraints. Similarly, a SIGCOMM 24' non-paper workshop [3] discussed five first-hand experiences of industry-academia collaboration, where each presenter interned at a large tech company (Cloudflare, Alibaba, and Prime Video). It also showed that 50% of SIGCOMM and NSDI papers now have at least one industry co-author.

In this paper, we also present industry-academia research collaboration models, but with the following differences. Unlike Google's editorial, which centers on a large tech company, our focus is on the perspective of a small tech company with a public mission. Additionally, our work differs from the SIGCOMM '24 workshop by not covering interns' experience working in the industry, but by showing multiple collaboration models beyond internships and academic papers publication.

Case study: we present a case study of *open research* collaboration between academic institutions and SIDN Labs [37], the research team of SIDN [63], a privately held small tech company with a public-interest mission. Our case study covers five different collaboration models (§3), which go beyond industry's proprietary datasets access (as in [41]) and beyond internships and papers (as in [3]). For each collaboration model, we provide deliverables which show how industry, academic partners, and the broader research and operator communities–including the SIGCOMM's –have benefited from the resulting open research.

Lastly, we also share lessons learned from building and running an in-house research team (§4). Our goal is to encourage other industry players to pursue open, collaborative research that serves academia, industry, and society alike.

2 BACKGROUND: SIDN AND SIDN LABS

SIDN is the registry and operator of the Netherlands' `.nl` country-code top-level domain (ccTLD). As such, it provides authoritative DNS services for `.nl` users [26] and maintains a database of all `.nl` domain names [60]. In 2024, SIDN's revenue was €25.6M, employing 99 staff members (90 FTE) [64].

Critical digital infrastructure: `.nl` domain name is of great importance in the Netherlands, being used by government, citizens, and businesses. As such, SIDN is classified by the Dutch government as part of the national digital critical infrastructure and is subject to the Network and Information Systems Security Act [5], which implements the EU's NIS1 directive [1].

SIDN Labs: SIDN Labs is the research team of SIDN. It carries out applied research to improve the security and robustness of the `.nl` ccTLD and the Internet infrastructure in the Netherlands and elsewhere. It has 15 members (11.5 FTEs) with complementary skills and perspectives, ranging from developers, operators, and researchers with master's and doctoral degrees. SIDN Labs engages in frequent collaboration with academic institutions, complementing its partnerships with industry [2, 24] and internal cooperation across SIDN teams. In this editorial, we focus exclusively on joint work with academia.

Research output: The output of our work consists of peer-reviewed papers, technical reports, open-source software [39], experimental network services (*e.g.,* secure time services [4]), and contributions to Internet standards [70], among others.

Funding: SIDN Labs' primary source of funding consists of 6% of SIDN's annual revenue, producing open research outputs as part of its public benefit mission.

3 COLLABORATION MODELS

Table 1 summarizes the collaboration models we have fostered with academia over the past thirteen years. We discuss if they need contracts, if they grant data access to SIDN's data, and require what kind of contributions (in-kind or in-cash), and examples of open deliverables produced by each model. Before diving into details, we discuss the importance of the human element in these collaborations.

3.1 Human element and incentives

Simply establishing a collaboration on paper does not ensure successful outcomes. Therefore, we suggest that both industry and academics carefully select staff who are genuinely committed to working together, otherwise it may not yield the expected results, as we have experienced in some cases.

Different incentives: the incentives for academic and industry may differ. Peer-reviewed publications are essential for academia, but for the industry, operation impact are typically more important while peer-reviewed publications may be secondary. Therefore, we recommend to clearly define the deliverables of each collaboration to manage expectations effectively, and make sure both sides remain committed and motivated.

3.2 Internships

Among the various forms of academic collaboration, student internships (Table 1) have proven to be the most straightforward to establish. Their costs are relatively low and the benefits they generate can pay off their costs. They last from 3 to 6 months. Setting up internships is relatively straightforward, as universities support internships in industry.

In this setup, the intern becomes part of our staff, and is supervised by one of our colleagues and by an academic at the university. This form of collaboration provides academics (in this case, the intern) access to our private datasets, subject to the signing of a non-disclosure agreement (NDA), given legal and privacy laws (*e.g.,* GDPR [19]).

Mutual benefits: As an industry partner, we profit from internships by having students work on well-defined research problems that we may lack the time or skills to address, thereby augmenting our research output. Several interns have then transitioned into staff positions. For academics, it offers students the opportunity to work with real-world datasets and problems, and gaining access to structured, production data, and acquire work experience in the industry. University supervisors are also involved in these internships, resulting in a master's thesis and, in some cases, a research paper as part of the output.

Takeaway: Internships are the low-hanging fruit of industry-academia collaboration, and can be a first step to build other types of collaboration. Industry investments are relatively low (mostly covering salary and supervision hours), and the benefits frequently surpass the investment. However, there are no guarantees, as students may not complete their projects as initially proposed, so they should not be on a project's critical path. Moreover, students can later transition into staff, particularly important in markets with a shortage of skilled workers.

3.3 Data sharing

Data sharing is another easy way to start collaborating with academics. It requires, however, setting up legal contracts and establishing data usage procedures. It can be done using aggregated data, anonymized, snapshots, and many other ways. (We share open aggregated datasets openly on our website [38]).

Case-by-case data sharing. Whenever we receive a data sharing request, we run a procedure internally to determine if we can (legally) share it. We require that the data be used solely for research and non-commercial purposes and that it benefits SIDN, the `.nl` domain, or society at large. We share data only with academic partners, and any sharing must conform to applicable privacy laws and ISO27001 certification [31].

We then evaluate every request individually. As an example, we share the `.nl` zone file daily with the University of Twente for the OpenIntel research project [58], which lists all delegated domain names in the zone and their DNS servers.

Mutual benefits: Data sharing enables academics to answer their research questions, and we make sure it benefits either SIDN and/or the community in general as a requisite to share the datasets.

Takeaway: We share datasets with academics to enable them to pursue their research questions; we find this is an easy way to collaborate.

3.4 Informal collaboration

Informal collaborations are the most prevalent form of academic partnership we utilize, where our staff work together with academics in research projects. They are also simple to establish, as they do not require contracts, NDAs or data sharing (Table 1).

Informal collaborations are not continuous; they usually have a common deliverable, such as a paper or prototype. Once a deliverable is completed, there is however a higher likelihood of researchers collaborating again on future studies. Researchers participate voluntarily in the projects in this model, funded by their own grants or employers. They are highly motivated to participate, given it is their own choice to do so. Among examples of informal collaboration, we list the series of 4 ACM IMC papers and one PAM paper we had on Authoritative DNS Server Engineering with USC/ISI and University of Twente [44, 46–48, 52], and our current work on Post-quantum DNSSEC [22, 34].

Mutual benefits: Both sides benefit from such collaborations. In our USC/ISI collaboration example, the recommendations we made for operators [42] were used by SIDN's operations team to improve

	Contract	Access to data	Cash contrib.	In-kind contrib.	Open deliverables
Internship	✓	✓	✓	✓	T [55], S [18]
Data sharing	✓	✓	✗	✓	D [58]
Informal	✗	✗	✗	✓	P [21, 45, 46, 48, 52, 53], S [67], R [42]
Secondment	✓	✗	✗	✓	P [43, 43, 50, 51, 54], S[57, 66], R [33, 49]
Funded research projects					
Self-funded	✓	✗	Optional	✓	P [44]
Externally funded	✓	✗	✗	✓	S [20], P [9, 10]
Commissioned	✓	✗	✗	✗	P [68], T [28, 35, 56]

Table 1: Industry-academia collaboration models. P (Papers), S (Software), R (RFCs), D (Datasets), Technical Reports (T)

their DNS services. For academics, it resulted in multiple tier-1 research papers of applied research and a better understanding of industry problems and requirements.

Takeaway: Informal collaborations are central to our research. They are efficient and straightforward to establish, and build bridges with academics. However, this form of collaboration often lacks continuity and funding, which can pose challenges for academics.

3.5 Secondment positions

While informal collaborations provide efficiency and speed, they often lack the sustained continuity and long-term objectives necessary for consistent industry and academic progress. To address these issues, we have established ongoing and long-term collaborations with selected academic institutions in the Netherlands through secondment positions, enabling a few of our staff to work at these institutions, one day per week (0.2 FTE).

Costs: For industry, secondment positions might seem expensive, as staff is allocated to work at the universities paid for by the industry partner. In practice, our secondments are structured so that our seconded staff works on projects that closely align with SIDN Labs' goals. This way, even though they are working elsewhere, they are contributing to both SIDN's mission and research goals.

Contracts: Our secondment contracts with universities have an intellectual property (IP) clause, which in most cases states that the universities own the intellectual property, but they license it to SIDN Labs unlimited. We are comfortable with this arrangement because we focus on open deliverables.

Current secondments: We have currently four colleagues (the authors of this paper) holding part-time secondment positions: two colleagues at the University of Twente, one at TU Delft, and one at the University of Amsterdam. They start as guests researchers at the universities, but are encouraged to climb the academic ladder at their hosting institutions – we have one full professor, one assistant professor, and two guest researchers.

External Ph.D. students: SIDN Labs staff members may use secondment constructs to pursue a Ph.D. degree at universities, provided that the regulations of the university permit and the work contributes to SIDN's mission. One of our team members completed his Ph.D. at the University of Twente under this arrangement researching DNSSEC [49], which is a core aspect of SIDN's operations.

Mutual benefits: This arrangement helps further skill development among our colleagues and keeps our research agenda updated as a result of knowledge exchange with the universities. It also enables our colleagues to maintain a continuous physical presence at various universities – spatial proximity has been previously associated with an increased propensity to collaborate [14, 62]. It also facilitates student internships, given our staff have contact with students.

Universities benefit from the contributions of the seconded staff through teaching, student supervision (in addition to in-house internships – §3.2) and joint-research proposals and projects.

Our seconded staff already has access to our datasets, so those can be leveraged in academic studies (*e.g.*, a phishing study[43]), as discussed in §3.3. Besides datasets, academics are exposed to our real-world problems, which can trigger further research.

Takeaway: Secondment positions are the most interactive and symbiotic form of academic collaboration we have and the only continuous form. In our experience, they have the potential to yield high benefits for the industry, academia, and the public at large.

3.6 Funded Research Projects

Industry and academia can collaborate in funded research projects (Table 1). Next we discuss the three types of funded research projects that we have been involved in.

3.6.1 *Self-co-funded research projects.*

We have co-funded several academic research projects (in-kind or in-cash) with other industry partners or funding agencies, which usually go to fund Ph.D. or post-doc positions. For example, we funded the TUCCR project [6] on cybersecurity with the University of Twente, together with other industry partners. Funding projects enables us to have a say over the research agenda, but we are mostly only involved in update meetings, and less in the day-to-day work of the researchers. As a result, we have limited control over the research process itself, unless one of our seconded staff (§3.5) is involved.

Expectations: Ph.D. students' main goal is to produce a dissertation, which means that the deliverables may not be at a technology readiness level (TRL) [23] that we can directly deploy. For example, a prototype may still require significant work before we can add it to our registration or DNS services. Also, Ph.D. students must focus on scientific output, not so much on industry impact. Similarly, post-docs also concentrate on scientific output to make a next step in their academic career.

Takeaway: Self-co-funded projects allow us more direct control over research agendas with academics, and enable academics to learn about industry's needs. These projects tend to be more costly and their deliverables will usually require extra work to reach higher TRLs. Industry-funded projects can be also an investment in training the new generation of engineers and researchers for society, which is how we look at it.

3.6.2 Externally-funded research projects. We also engage in externally-funded research projects that bring together industry and academia, usually by governmental funding agencies. In these types of projects, our work gets partially funded too.

Mutual benefits: Solutions developed in projects like these can benefit multiple parties. For example, we joined the EU Horizon CONCORDIA [16] research project, a €23M project involving 56 academic and industry partners that included funding for industry partners such as us. We worked there to develop a DDoS clearing house [20], an open-source tool for generating and sharing DDoS fingerprints. All partners benefited from mutual work, and the projected allowed us as a research team to expand our research network, into other sectors and other industry and academic partners.

For us in the industry, government-funded projects provide visibility, allow us to increase our output, and offer the opportunity to work and learn from new partners. Academics also benefit from domain experts available in the industry, who have deep knowledge in the deployment and operations of Internet services. Costwise, government-funded projects are a way to obtain external co-funding for industry research projects.

3.6.3 Commissioned research projects. As an industry lab, we also engage in commissioned research projects , with other industry partners, operators [36], and other universities. In such projects, a funding agency typically opens a call for proposals for specific research questions they need to have answered, and industry labs can work as contractors, together with academics. We carried out multiple of such projects with ICANN [12, 28, 35, 56] and for the Government of the Netherlands [68].

The main advantage of commissioned projects is that they allow us to apply our expertise in other projects. We take on these projects to support the community and to demonstrate the relevance of our work, not to generate revenue. The commissioning organization, however, defines the project scope and owns the resulting outputs.

4 INDUSTRY RESEARCH TEAMS

There have been multiple studies [7, 8, 15, 25, 61] showing the benefits of research collaborations between academia and industry, such as knowledge transfer, development of new products, processes, and economical and community benefits and services. Below we summarize advantages and lessons learned based our direct experience in collaboration with academia.

4.1 Advantages for Industry

Reducing dark data and improving services: Organizations collect, process, and store data during regular business activities but often find it difficult to maximize the use for any meaningful analysis or decision-making – which qualifies this data as "dark data" [17]. Collaborations with academia can help to draw more value from such datasets, for instance with the help of students [11, 27]. Lessons learned from these datasets' analysis can be used to improve processes and services, as, for example, we did with `.nl` authoritative servers [42].

Scrutinizing operator's decisions. As a DNS operator, SIDN's engineers focus on running large distributed DNS services and do not have either the time or experience to systematically carry out evaluations of various DNS configurations parameters. As the saying goes: "if it ain't broke, don't fix it".

Researchers from Labs, however, freed from the duty of running production networks, can carry out studies to asses and further improve choices made by the engineers. For example, what should be the right time-to-live value for DNS records, given they range from 0s to 68 years [40]? We carried out two studies to answer this question [46, 48]. Also, we studied how recursive resolvers reach our authoritative name servers which helped our DNS engineers to optimize the setup [42, 52].

Policy review: The impact of operator's policies is an important subject for academic scrutiny. For instance, we collaborated with two other registries (Belgium's `.be` and Ireland's `.ie`) and four universities to conduct a comparative analysis of phishing attacks and their mitigation across the three ccTLDs [43]. Based on the insights from this study, SIDN is currently reevaluating the phishing mitigation policies for .nl.

Contributing to the public good: Open research means that findings are publicly shared. Consequently, any potential lessons and discoveries can benefit the public – for instance, through a RFC for large DNS operators [42] or open-source software for DNS and Internet traffic analysis [65]. For public-benefit organizations such as SIDN, this serves as a strong incentive for academic collaboration.

4.2 Advantages for Academia

Academia and academics also benefit from working with industry. We list the benefits we observed from our experience.

Data access: Industry frequently possess unique datasets that can be leveraged for academic research. In an era where "data is the new oil", access to these datasets can yield significant insights and contributions. The types of data held by industry are often unknown to academics, and it is through direct data exploration of these datasets that their potential can be unearthed, as in the case of OpenIntel and multiple studies that have used it [13, 69].

Real-world problems: Industry partners provide real-world services and products, and face practical problems that need solutions. Ph.D. students, M.Sc. students, and research projects can use these problems to further increase the relevance of their work for industry and society, and have a direct impact on operations.

Domain experts: Industry's experts are responsible for operating real-world systems, such as SIDN's DNS engineers that run `.nl`. These domain experts have accumulated substantial knowledge over time, which can be a valuable resource for academic research projects. For comparison, a typical computer science undergraduate spends approximately a couple of hours studying DNS. In contrast, there are 297 RFCs totaling 2082 pages of DNS documentation [30] in IETF documents – the primary learning resource for operators. As such, industry domain experts have a deep knowledge and hands-on experience that can be leveraged by academics, and they often have research problems sought by academics.

Funding: Industry partners can also sponsor academic research, either by sponsoring projects or in-kind contributions (§3.6).

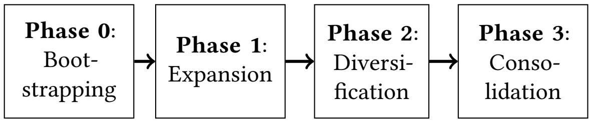

Figure 1: Research Team Evolution

Networks: Industry can offer academics access to industry networks, potentially paving the way for additional collaboration opportunities, such as with RIPE, IETF, and CENTR in our case. For example, bachelor and masters students supervised by SIDN Labs and TU Delft delivered presentations on NTP at the RIPE91 [32, 59], and at ICANN84 [29] meetings, both in October 2025.

4.3 Building an Industry research team

For small companies aiming to establish research teams in collaboration with academia, we outline the key phases that have shaped our own partnerships (Figure 1). Interested readers can find additional details in [24].

Phase 0: Bootstrapping. After the board approved the establishment of an internal research team, we initially focused on in-cash sponsorship of Ph.D. positions at Dutch universities to foster ties with academia (§3.6.1).

Phase 1: Expansion. In this phase, we expanded our research team with members from SIDN's internal development and operations teams, and built a dedicated research infrastructure of networks and systems, while sponsoring external academic projects.

Phase 2: Diversification. We further grew the team with new members, including additional academically trained researchers. This enabled our first informal academic collaborations (§3.4), while maintaining the in-cash sponsorship model from Phase 1. The increased internal capacity provided greater influence over research topics and over the transfer of results into production and standardization. During this phase, we also initiated other collaboration models, such as internships and data sharing.

Phase 3: Consolidation. We further grew the team, taking up topics such as machine learning for domain name and Internet security. We found it more beneficial and economical to hire our own researchers rather than focusing on sponsoring temporary Ph.D. and postdoc projects. This approach allowed our researchers to maintain collaborations with academia, focusing on in-kind contributions instead of in-cash funding. Also, this is when we set up secondment positions (§3.5). This phase allowed to keep talent in-house and further increasing control over the research directions.

5 OUTLOOK

We presented five models for open research collaboration between industry and academia, using a small tech company as a case study. Based on over thirteen years of experience, we discussed practical insights for how small industry players can build research capacity and engage with academic partners. We hope these lessons are useful to others in establishing academic collaborations to the benefit of the industry, academia, and the community at large.

Acknowledgments. We thank our collaborators from academia and industry for their contributions over the past thirteen years.

[1] 2022. Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union, amending Regulation (EU) No 910/2014 and Directive (EU) 2018/1972 (NIS2 Directive). (2022). https://eur-lex.europa.eu/legal-content/en/TXT/?uri=CELEX%3A32022L2555 Accessed: 2025-02-24.

[2] 2024. About SIDN Labs. (2024). https://www.sidnlabs.nl/en/about-sidnlabs Accessed: 2024-11-28.

[3] 2024. SIGCOMM'24: The Role of Industry and Academia in Network Research. (2024). https://conferences.sigcomm.org/sigcomm/2024/nonpaper/#industry-academia-role Video recording available at https://www.youtube.com/watch?v=0T9roF0mg7g.

[4] 2024. TimeNL - Public NTP Service. (2024). https://time.nl/index_en.html

[5] 2024. Wet beveiliging netwerk- en informatiesystemen (*In Dutch*). (2024). https://wetten.overheid.nl/BWBR0041515/2024-10-01 Accessed: 2025-01-23.

[6] 2025. Twente University Centre for Cybersecurity Research. (2025). https://www.utwente.nl/en/digital-society/research/cybersecurity_tuccr/ Accessed: 2025-02-14.

[7] Roberta Apa, Valentina De Marchi, Roberto Grandinetti, and Silvia Rita Sedita. 2021. University-SME collaboration and innovation performance: the role of informal relationships and absorptive capacity. *The Journal of Technology Transfer* 46, 4 (2021), 961–988. https://doi.org/10.1007/s10961-020-09802-9

[8] João Barbosa, Gabriela Fernandes, and Anabela Tereso. 2023. Benefits of University-Industry R&D Collaborations: A Systematic Literature Review. In *Innovations in Industrial Engineering II*, José Machado, Filomena Soares, Justyna Trojanowska, Vitalii Ivanov, Katarzyna Antosz, Yi Ren, Vijaya Kumar Manupati, and Alejandro Pereira (Eds.). Springer International Publishing, Cham, 257–280.

[9] Antonio Battipaglia, Leonardo Boldrini, Ralph Koning, and Paola Grosso. 2023. Evaluation of SCION for User-driven Path Control: a Usability Study. In *Proceedings of the SC '23 Workshops of the International Conference on High Performance Computing, Network, Storage, and Analysis (SC-W '23)*. Association for Computing Machinery, New York, NY, USA, 785–794. https://doi.org/10.1145/3624062.3624592

[10] Rodrigo Bazo, Leonardo Boldrini, Cristian Hesselman, and Paola Grosso. 2021. Increasing the Transparency, Accountability and Controllability of multi-domain networks with the UPIN framework. In *Proceedings of the ACM SIGCOMM 2021 Workshop on Technologies, Applications, and Uses of a Responsible Internet (TAURIN'21)*. Association for Computing Machinery, New York, NY, USA, 8–13. https://doi.org/10.1145/3472951.3473506

[11] Rushvanth Bhaskar. 2022. *A Day in the Life of NTP: Analysis of NTPPool Traffic*. Master's thesis. University of Twente and SIDN Labs, Enschede and Arnhem, The Netherlands. Master's thesis.

[12] CDAR. [n. d.]. Continuous Data-driven Analysis of Root Stability (CDAR). ([n. d.]). https://cdar.nl/ Accessed: 2024-11-22.

[13] Taejoong Chung, Roland van Rijswijk-Deij, David Choffnes, Dave Levin, Bruce M Maggs, Alan Mislove, and Christo Wilson. 2017. Understanding the role of registrars in DNSSEC deployment. In *Proceedings of the 2017 Internet Measurement Conference*. 369–383.

[14] Matthew Claudel, Emanuele Massaro, Paolo Santi, Fiona Murray, and Carlo Ratti. 2017. An exploration of collaborative scientific production at MIT through spatial organization and institutional affiliation. *PloS one* 12, 6 (2017), e0179334.

[15] European Commission. 2024. The Future of European Competitiveness: A Competitiveness Strategy for Europe. (2024). https://commission.europa.eu/document/download/97e481fd-2dc3-412d-be4c-f152a8232961_en Accessed: 2024-12-03.

[16] CONCORDIA Consortium. 2024. CONCORDIA: A Cybersecurity Competence Network. https://www.concordia-h2020.eu/. (2024). Accessed: 2024-07-08.

[17] Angelo Corallo, Anna Maria Crespino, Vito Del Vecchio, Mariangela Lazoi, and Manuela Marra. 2023. Understanding and Defining Dark Data for the Manufacturing Industry. *IEEE Transactions on Engineering Management* 70, 2 (2023), 700–712. https://doi.org/10.1109/TEM.2021.3051981

[18] Doesburg, Dirk. 2025. PQC for the RPKI. https://github.com/SIDN/pqc-rpki. (July 2025).

[19] European Parliament and Council of the European Union. [n. d.]. Regulation (EU) 2016/679 of the European Parliament and of the Council. ([n. d.]). https://data.europa.eu/eli/reg/2016/679/oj

[20] EU's Horizon 2020 Concordia project. 2022. DDoS Clearing House. https://github.com/ddos-clearing-house/. (2022). A modular application for combating DDoS attacks. Co-funded by EU Horizon 2020 program (grant number 830927).

[21] Ginevra Fabrizio, Ralph Koning, Elmer Lastdrager, Caspar Schutijser, Anna Sperotto, and Roland van Rijswijk-Deij. 2025. PQC for DNSSEC: a format size analysis on Falcon signatures. In *Proceedings of the 2025 Applied Networking Research Workshop (ANRW '25)*. Association for Computing Machinery, New York, NY, USA, 143–149. https://doi.org/10.1145/3744200.3744767

[22] Andrew Fregly, Roland van Rijswijk-Deij, Moritz Müller, Peter Thomassen, Caspar Schutijser, and Taejoong Chung. 2024. *Research Agenda for a Post-Quantum DNSSEC*. Internet-Draft draft-fregly-research-agenda-for-pqc-dnssec-02. Internet Engineering Task Force. https://datatracker.ietf.org/doc/draft-fregly-research-a

genda-for-pqc-dnssec/02/ Work in Progress.
[23] Mihály Héder. 2017. From NASA to EU: the evolution of the TRL scale in Public Sector Innovation. *The Innovation Journal* 22, 2 (2017), 1–23.
[24] Cristian Hesselman. 2021. Ten years of SIDN Labs. (2021). https://www.sidn.nl/en/news-and-blogs/10-years-of-sidn-fund-for-a-strong-open-and-free-internet
[25] Frederic Hilkenmeier, Christian Fechtelpeter, and Julian Decius. 2021. How to foster innovation in SMEs: evidence of the effectiveness of a project-based technology transfer approach. *The Journal of Technology Transfer* (2021). https://doi.org/10.1007/s10961-021-09913-x
[26] Paul E. Hoffman and Kazunori Fujiwara. 2024. *DNS Terminology.* RFC 9499. https://doi.org/10.17487/RFC9499
[27] Pascal Huppert. 2024. *Identifying DNS Scanners from a TLD Perspective.* Master's thesis. University of Münster and SIDN Labs, Münster, Germany and Arnhem, The Netherlands. Master's thesis.
[28] ICANN. 2024. *Reduced Risk of Redirected Query Traffic with Signed Root Name Server Data.* Technical Report. https://www.icann.org/en/system/files/files/reduced-risk-redirected-query-traffic-signed-root-name-server-data-22may24-en.pdf
[29] ICANN 84 Conference. 2025. ccNSO Tech Day (1 of 3). (2025). https://icann84.sched.com/event/29QJv/ccnso-tech-day-1-of-3 Accessed: 2025-11-13.
[30] IETF. 2024. Herding the DNS Camel. [Online]. Available: https://powerdns.org/dns-camel/. (2024). Accessed: Jul. 3, 2024.
[31] International Organization for Standardization. 2022. *ISO/IEC 27001:2022 - Information Security Management Systems Requirements.* ISO, Geneva, Switzerland. Accessed: 2024-12-02.
[32] Shreyas Konjerla. 2025. Are NTP Clients Always Right? (2025). https://ripe91.ripe.net/programme/meeting-plan/sessions/17/RNBMHH/ Presented at RIPE 91: Tuesday, 21 October 2025, 11:30 (UTC+03:00), Main Room. Accessed: 2025-11-12.
[33] S. Konjerla, G. C. M. Moura, G. Smaragdakis, and T. Dittrich. 2025. *Are NTP Clients Always Right? Evaluating NTP Clients under Normal and Attack Scenarios.* Technical Report SIDN Labs Technical Report 2025-10-16. SIDN Labs and TU Delft, Arnhem, The Netherlands. https://www.sidnlabs.nl/downloads/7M1bz1otGu0D7hqr0hRT2c/3bc09536c3bc87f63b80d66944abc1d9/ntp-clients-tech_report-20251016.pdf Updated October 30, 2025.
[34] SIDN Labs. 2023. PATAD: A Testbed for Evaluating Quantum-Safe Cryptography in DNSSEC. (2023). https://patad.sidnlabs.nl/ Accessed: 2025-01-21.
[35] SIDN Labs. 2023. *RSSAC028 Implementation Study Report.* Technical Report. https://www.sidnlabs.nl/downloads/xPHs3U5y8LNunHe9wCarI/32045255e07df7e5e1a35b98d8be4deb/rssac028-implementation-study-report-27sep23-en.pdf
[36] SIDN Labs. 2024. DNS Belgium and SIDN Collaborate on ML Project Aimed at Detecting Suspect Domain Name Registrations. (December 2024). https://sidnlabs.nl/en/news-and-blogs/dns-belgium-and-sidn-collaborate-on-ml-project-aimed-at-detecting-suspect-domain-name-registrations
[37] SIDN Labs. 2024. SIDN Labs. (July 2024). https://www.sidnlabs.nl/en
[38] SIDN Labs. 2024. SIDN Labs Statistics. (2024). https://stats.sidnlabs.nl/en/ Accessed: 2024-12-05.
[39] SIDN Labs. 2024. Tools and Measurements. (2024). https://www.sidnlabs.nl/en/tools-and-measurements Accessed: 2024-12-03.
[40] Paul Mockapetris. 1987. *Domain names - concepts and facilities.* RFC 1034. IETF. http://tools.ietf.org/rfc/rfc1034.txt
[41] Jeffrey C. Mogul, Priya Mahadevan, Christophe Diot, John Wilkes, Phillipa Gill, and Amin Vahdat. 2021. Data-driven networking research: models for academic collaboration with industry (a Google point of view). *SIGCOMM Comput. Commun. Rev.* 51, 4 (dec 2021), 47–49. https://doi.org/10.1145/3503954.3503960
[42] G. Moura, W. Hardaker, J. Heidemann, and M. Davids. 2022. *Considerations for Large Authoritative DNS Server Operators.* RFC 9199. IETF. http://tools.ietf.org/rfc/rfc9199.txt
[43] Giovane C. M. Moura, Thomas Daniels, Maarten Bosteels, Sebastian Castro, Moritz Müller, Thymen Wabeke, Thijs van den Hout, Maciej Korczyński, and Georgios Smaragdakis. 2024. Characterizing and Mitigating Phishing Attacks at ccTLD Scale. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security (CCS '24).* Association for Computing Machinery, New York, NY, USA, 2147–2161. https://doi.org/10.1145/3658644.3690192
[44] Giovane C. M. Moura, Ricardo de O. Schmidt, John Heidemann, Wouter B. de Vries, Moritz Müller, Lan Wei, and Christian Hesselman. 2016. Anycast vs. DDoS: Evaluating the November 2015 Root DNS Event. In *Proceedings of the ACM Internet Measurement Conference.* ACM, Santa Monica, California, USA, 255–270. https://doi.org/10.1145/2987443.2987446
[45] Giovane C. M. Moura and John Heidemann. 2023. Vulnerability Disclosure Considered Stressful. *SIGCOMM Comput. Commun. Rev.* 53, 2 (jul 2023), 2–10. https://doi.org/10.1145/3610381.3610383
[46] Giovane C. M. Moura, John Heidemann, Ricardo de O. Schmidt, and Wes Hardaker. 2019. Cache Me If You Can: Effects of DNS Time-to-Live. In *Proceedings of the ACM Internet Measurement Conference.* ACM, Amsterdam, the Netherlands, 101–115. https://doi.org/10.1145/3355369.3355568
[47] Giovane C. M. Moura, John Heidemann, Wes Hardaker, Pithayuth Charnsethikul, Jeroen Bulten, João M. Ceron, and Cristian Hesselman. 2022. Old but Gold: Prospecting TCP to Engineer and Live Monitor DNS Anycast. In *Proceedings of the Passive and Active Measurement Workshop.* Springer, virtual, to appear. https://doi.org/tbd Awarded best paper.
[48] Giovane C. M. Moura, John Heidemann, Moritz Müller, Ricardo de O. Schmidt, and Marco Davids. 2018. When the Dike Breaks: Dissecting DNS Defenses During DDoS. In *Proceedings of the ACM Internet Measurement Conference.* ACM, Boston, MA, USA, 8–21. https://doi.org/10.1145/3278532.3278534
[49] Moritz Müller. 2021. *Making DNSSEC Future Proof.* PhD Thesis - Research UT, graduation UT. University of Twente, Netherlands. https://doi.org/10.3990/1.9789036551816
[50] Moritz Müller, Taejoong Chung, Alan Mislove, and Roland van Rijswijk-Deij. 2019. Rolling With Confidence: Managing the Complexity of DNSSEC Operations. *IEEE Transactions on Network and Service Management* 16, 3 (2019), 1199–1211. https://doi.org/10.1109/TNSM.2019.2916176
[51] Moritz Müller, Jins de Jong, Maran van Heesch, Benno Overeinder, and Roland van Rijswijk-Deij. 2020. Retrofitting post-quantum cryptography in Internet protocols: a case study of DNSSEC. *SIGCOMM Comput. Commun. Rev.* 50, 4 (Oct. 2020), 49–57. https://doi.org/10.1145/3431832.3431838
[52] Moritz Müller, Giovane C. M. Moura, Ricardo de O. Schmidt, and John Heidemann. 2017. Recursives in the Wild: Engineering Authoritative DNS Servers. In *Proceedings of the ACM Internet Measurement Conference.* ACM, London, UK, 489–495. https://doi.org/10.1145/3131365.3131366
[53] Moritz Müller, Matthew Thomas, Duane Wessels, Wes Hardaker, Taejoong Chung, Willem Toorop, and Roland van Rijswijk-Deij. 2019. Roll, Roll, Roll your Root: A Comprehensive Analysis of the First Ever DNSSEC Root KSK Rollover. In *Proceedings of the Internet Measurement Conference (IMC '19).* Association for Computing Machinery, New York, NY, USA, 1–14. https://doi.org/10.1145/3355369.3355570
[54] Moritz Müller, Willem Toorop, Taejoong Chung, Jelte Jansen, and Roland van Rijswijk-Deij. 2020. The Reality of Algorithm Agility: Studying the DNSSEC Algorithm Life-Cycle. In *Proceedings of the ACM Internet Measurement Conference (IMC '20).* Association for Computing Machinery, New York, NY, USA, 295–308. https://doi.org/10.1145/3419394.3423638
[55] Multiple authors. 2025. Bachelor and master theses published at SIDN Labs. https://www.sidnlabs.nl/en/publications#theses. (July 2025).
[56] Moritz Müller, Jelte Jansen, Marco Davids, and Willem Toorop. 2022. *DNSSEC Deployment Metrics Research.* Technical Report. SIDN Labs and NLnet Labs. https://www.icann.org/en/system/files/files/dnssec-deployment-metrics-research-08aug22-en.pdf
[57] NTPinfo Developers. 2025. NTPinfo: An open-source tool for evaluating NTP server accuracy with also using RIPE Atlas probes. https://github.com/NTPinfo/NTPinfo. (2025). Accessed: 2025-07-09.
[58] University of Twente. 2024. OpenINTEL. (2024). https://openintel.nl/
[59] Calin Olaru, Horia-Andrei Botezatu, Mihai Nicolae, and Serban Orza. 2025. NTPinfo: Application for Measuring NTP Servers Accuracy and More. (2025). https://ripe91.ripe.net/programme/meeting-plan/sessions/17/8WYZGT/ Presented at RIPE 91, Tuesday, 21 October 2025. Accessed: 2025-11-12.
[60] J. Postel. 1994. *Domain Name System Structure and Delegation.* RFC 1591. IETF. http://tools.ietf.org/rfc/rfc1591.txt
[61] Sergio Rico, Felix Dobslaw, and Lena-Maria Öberg. 2024. Starting Collaborations Between SMEs and Researchers in Software Engineering. In *Software Business*, Sami Hyrynsalmi, Jürgen Münch, Kari Smolander, and Jorge Melegati (Eds.). Springer Nature Switzerland, Cham, 222–230.
[62] Arianna Salazar Miranda and Matthew Claudel. 2021. Spatial proximity matters: A study on collaboration. *Plos one* 16, 12 (2021), e0259965.
[63] SIDN. 2024. SIDN. (July 2024). https://www.sidn.nl/en
[64] SIDN. 2025. SIDN Annual Report 2024: "Global interconnection is the internet's strength". (14 May 2025). https://www.sidn.nl/en/about-sidn/annual-report-2024 Accessed: 2025-08-25.
[65] SIDN Labs. 2024. ENTRADA - DNS Big Data Analytics. (Jan. 2024). https://entrada.sidnlabs.nl/.
[66] SIDN Labs. 2025. NTP Info. https://ntpinfo.sidnlabs.nl. (2025). https://ntpinfo.sidnlabs.nl Accessed: 2025-07-09.
[67] SIDN Labs and InternetNZ and USC/ISI. 2021. CycleHunter: Detecting cyclic dependencies in DNS zone files. https://github.com/SIDN/CycleHunter. (6 May 2021). https://github.com/SIDN/CycleHunter Python tool, BSD-2-Clause license.
[68] Raffaele Sommese, Mattijs Jonker, Jeroen van der Ham, and Giovane C. M. Moura. 2022. Assessing e-Government DNS Resilience. In *2022 18th International Conference on Network and Service Management (CNSM).* 118–126. https://doi.org/10.23919/CNSM55787.2022.9965155
[69] Olivier Van Der Toorn, Roland van Rijswijk-Deij, Tobias Fiebig, Martina Lindorfer, and Anna Sperotto. 2020. TXTing 101: finding security issues in the long tail of DNS TXT records. In *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW).* IEEE, 544–549.
[70] Maarten Wullink and Paweł Kowalik. 2025. *RESTful Provisioning Protocol (RPP).* Internet-Draft draft-wullink-rpp-core-03. Internet Engineering Task Force. https://datatracker.ietf.org/doc/draft-wullink-rpp-core/03/ Work in Progress.